RHCE Red Hat Certified Engineer Linux Study Guide (Exam RH302), Fifth Edition

byMichael Jang

McGraw-Hill/Osborne 2007 (896 pages)
ISBN:9780072264548

With hundreds of practice questions and hands-on exercises, this authoritative guide covers what you need to know--and shows you how to prepare--for the challenging RHCE exam (RH302).

Table of Contents

Back Cover

The Best Fully Integrated Study System Available

With hundreds of practice questions and hands-on exercises, *RHCE Red Hat Certified Engineer Linux Study Guide, Fifth Edition* covers what you need to know--and shows you how to prepare--for this challenging exam.

- 100% complete coverage of all objectives for exam RH302
- Exam Readiness Checklist at the front of the book--you're ready for the exam when all objectives on the list are checked off
- Inside the Exam sections in every chapter highlight key exam topics covered
- Real-world exercises modeled after hands-on exam scenarios
- Two complete lab-based exams simulate the format, tone, topics, and difficulty of the real exam

Covers all RH302 exam topics, including:

- Hardware installation and configuration
- The boot process
- Linux filesystem administration
- Package management and Kickstart
- User and group administration
- System administration tools
- Kernel services and configuration
- Apache and Squid
- Network file sharing services (NFS, FTP, and Samba)
- Domain Name System (DNS)
- E-mail (servers and clients)
- Extended Internet Services Daemon (xinetd), the Secure package, and DHCP
- The X Window System
- Firewalls, SELinux, and troubleshooting

About the Author

Michael Jang (RHCE, LPIC-1, LCP, Linux+, MCP) is currently a full-time writer, specializing in operating systems and networks. His experience with computers goes back to the days of jumbled punch cards. He has written other books on Linux certification, including *Linux+ Exam Cram* and *Sair GNU/Linux Installation and Configuration Exam Cram*. His other Linux books include *Linux Annoyances for Geeks, Linux Patch Management*, and *Mastering Fedora Core Linux 5*. He has also written or contributed to books on Microsoft operating systems, including *MCSE Guide to Microsoft Windows 98 and Mastering Windows XP Professional, Second Edition*.

RHCE Red Hat Certified Engineer Linux Study Guide (Exam RH302), Fifth Edition

Michael Jang

New York, Chicago, San Francisco, Lisbon, London, Madrid, Mexico City, Milan, New Delhi, San Juan, Seoul, Singapore, Sydney Toronto

Cataloging-in-Publication Data is on file with the Library of Congress

McGraw-Hill books are available at special quantity discounts to use as premiums and sales promotions, or for use in corporate training programs. For more information, please write to the Director of Special Sales, Professional Publishing, McGraw-Hill, Two Penn Plaza, New York, NY 10121-2298. Or contact your local bookstore.

1234567890 CUS CUS 01987

ISBN-13: 978-0-07-226454-8

ISBN-10:
0-07-226454-3

Sponsoring Editor
Timothy Green

Editorial Supervisor
Janet Walden

Project Editor
LeeAnn Pickrell

Acquisitions Coordinator
Jennifer Housh

Technical Editor
Elizabeth Zinkann

Copy Editor
Lisa Theobald

Proofreader
Paul Tyler

Indexer
Rebecca Plunket

Production Supervisor
Jim Kussow

Composition
Apollo Publishing Services

Illustration
Apollo Publishing Services

Art Director, Cover
Jeff Weeks

Cover Designer
Pattie Lee

For the young widows and widowers, may they find the courage to face their fears, to navigate their way through the pain, and to find hope for a brighter future.

About the Contributors

Author

Michael Jang (RHCE, LPIC-1, LCP, Linux+, MCP) is currently a full-time writer, specializing in operating systems and networks. His experience with computers goes back to the days of jumbled punch cards. He has written other books on Linux certification, including *Linux+ Exam Cram* and *Sair GNU/Linux Installation and Configuration Exam Cram*. His other Linux books include *Linux Annoyances for Geeks*, *Linux Patch Management*, and *Mastering Fedora Core Linux 5*. He has also written or contributed to books on Microsoft operating systems, including *MCSE Guide to Microsoft Windows 98* and *Mastering Windows XP Professional, Second Edition*.

Technical Editor

Elizabeth Zinkann is a logical Linux catalyst, a freelance technical editor, and an independent computer consultant.

She was a contributing editor and review columnist for *Sys Admin Magazine* for ten years. As an editor, some of her projects have included *Mastering Fedora Core Linux 5*, *Ending Spam*, *Linux Patch Management*, and *Write Portable Code*. In a former life, she also programmed communications features, including ISDN at AT&T Network Systems.

Acknowledgments

I personally would like to thank the following people:

- **Nancy E. Cropley, R.N. (d. 2002)** It's now been over five years since you've left this world, but I continue to hold your spirit in my heart, and I hope you can still see the joy of the world through my eyes. You are my hero, even today. I hope you can see how happy I am with Donna, but I wish I could still be with you. I will always miss you.

  As a political activist, you fought for what you believed in: social justice, peace, and universal healthcare. You were never afraid to go to jail to support your beliefs. Your example is helping me find a backbone for life.

  As a nurse for the homeless, you helped so many who are less fortunate. You worked tirelessly in the clinics, in the shelters, and on the streets. Your efforts eased the pain of so many people. And you saved lives.

  As an Internet entrepreneur, you showed me how to be happy pursuing a life working from home. You made it possible for me to have the freedom to be, instead of getting stuck in the corporate world.

  Nancy, you were my partner, my lover, my soul mate. You helped me find joy in this world. I take your lessons with me. I thank you for the best seven years of my life.

- All the incredibly hard-working folks at McGraw-Hill: Tim Green, Jennifer Housh, LeeAnn Pickrell, Lisa Theobald, Paul Tyler, and Rebecca Plunket for their help in launching a great series and being solid team players.

Preface

Overview

Linux is thriving. Red Hat is at the forefront of the Linux revolution. And Red Hat Certified Engineers and Technicians are making it happen.

Even in the current economic recovery, business, education, and governments are cost conscious. They want control of their operating systems. Linux-even Red Hat Enterprise Linux-saves money. The open source nature of Linux allows users to control and customize their operating systems. While there is a price associated with Red Hat Enterprise Linux (RHEL), the cost includes updates and support. Now with Xen, it's possible to set up a cluster of virtual, independent installations of RHEL (and other operating systems) on a single physical computer. As I describe shortly, there are freely available "rebuilds" of RHEL that you can get without support from Red Hat, with features identical for most administrators.

On the Job

A "rebuild" is software that is built by a third party from the same source code as the original "build." On the other hand, a "clone" is built from different source code.

As this book is going to print, the New York Stock Exchange has just announced that it's moving to Linux. Major corporations, from Home Depot to Toyota, and governments such as Brazil, the Republic of Korea, and Switzerland have made the switch to Linux. When faced with a Microsoft audit for licenses, the Portland, Oregon, school system switched to Linux. Major movie studios such as Disney and Dreamworks use Linux to create the latest motion pictures. IBM has invested billions in Linux -and constantly features Linux in its advertising. HP has reported 2.5 billion dollars in Linux-related revenue in 2003, and it's still growing today (2007). Even though Linux is freely downloadable, Wall Street Technology just reported that Linux server revenue in 2006 was about 7 billion dollars, 1/3rd that of Microsoft (up from 1/4th in 2004), and is still gaining market share. Is Microsoft Vista motivating business to look more closely at Linux?

With the One Laptop Per Child (OLPC) initiative, a streamlined version of Fedora Core 6 will be placed in front of tens (or possibly hundreds) of millions of students worldwide. These students will learn Linux first. And Red Hat Enterprise Linux 5 is based on Fedora Core 6.

Security is another reason to move toward Linux. The U.S. National Security Agency has developed its own version of the Linux kernel to provide context-based security; RHEL has incorporated many of these improvements.

While there are Linux distributions available from a number of companies, Red Hat is far and away the market leader. Novell's acquisition of SUSE hasn't made a dent. Based on 2006 sales, Red Hat has apparently shrugged off the challenge of Oracle Linux (which is another "rebuild" of Red Hat Enterprise Linux). Incidentally, the RHCE was named #1 in CertCities.com's list of hottest certifications for 2006. Therefore, the RHCE provides the most credibility to you as a Linux professional.

The RHCT and RHCE exams are difficult. Available historical data suggests that less than 50 percent of first-time candidates pass the RHCE exam. But do not be intimidated. While there are no guarantees, this book can help you prepare for and pass the Red Hat Certified Technician and Red Hat Certified Engineer exams. And these same skills can help you in your career as a Linux administrator. Just remember, this book is not intended to be a substitute for Red Hat prep courses that I describe shortly.

To study for this exam, you should have a network of at least two Linux or Unix computers. (It's acceptable if these computers are on virtual machines such as VMware or Xen.) You need to install RHEL on at least one of these computers. That will allow you to configure Linux and test the results. After configuring a service, especially a network service, it's important to be able to check your work from another computer.

Getting Red Hat Enterprise Linux

The Red Hat exams are based on your knowledge of *Red Hat* Enterprise Linux. When you take the RHCE exam, it'll be on a "standard" PC with Intel 32-bit (or compatible) personal computers. The CPU should have a speed of at least 700MHz, and the PC should have at least 256MB of RAM. As Red Hat Network updates are not explicitly listed as a requirement in the Red Hat Exam Prep guide, a "trial" subscription or a rebuild distribution is probably sufficient. If you want a full subscription, which can help you test features associated with the Red Hat Network, the price depends on your hardware and the amount of support you need. I've emphasized *Red Hat* solely to focus on distributions that use Red Hat source code, including the "rebuilds" described in this section (and more).

With Red Hat Enterprise Linux 5, Red Hat has modified its offerings into two categories:

- RHEL Server includes varying levels of support for entry-level to high-end and mission-critical systems.
    - The RHEL Server Advanced Platform supports unlimited virtualized guests, virtualized storage, high-availability clustering and failover, with support for more than two CPUs.
    - RHEL Server subscriptions are available for IBM System Z mainframe systems on a per-processor basis.
    - RHEL Server subscriptions are also available for High Performance Computing clusters.
- RHEL Desktop includes varying levels of support suitable for desktop computers and workstations. There are different options available for systems with one or more CPUs.

If you want to prepare for the RHCE exam with the official RHEL 5 server operating system, trial subscriptions are available (www.redhat.com/en_us/USA/home/developer/trial/). While they only support updates for 30 days, updates can also be tested using the mirror repositories associated with rebuild distributions. And you can download the same operating system (for the trial period) from the same sources as paying Red Hat users.

But you don't have to pay for the operating system or settle for a "trial subscription" to prepare for the RHCE exam. There are a wide variety of efforts to create "rebuilds" of Red Hat Enterprise Linux. The source code for almost all RHEL RPM packages is released under the Linux General Public License (GPL) or related licenses. This gives anyone the right to build Red Hat Enterprise Linux from the Red Hat released source code.

The source code is released in Source RPM package format, which means the RPM packages can be built using the **rpm** commands described in Chapter 5. The developers behind rebuild distributions have all revised the source code to remove Red Hat trademarks. Most, like CentOS-5, are freely available; others, like Oracle Linux, require a subscription.

On the Job

Oracle Linux has tried to undercut Red Hat by developing their own rebuild of Red Hat Enterprise Linux. Their subscriptions cost less at what I presume are similar support levels. As I have not tried Oracle Linux, I do not know if you get the same level of knowledge that you would get from Red Hat engineers.

You can select and download the rebuild that most closely meets your needs. I have tried several of the rebuilds, including those developed by Community Enterprise Linux (CentOS), Scientific Linux, and Lineox. All have proven reliable. In fact, they are so popular, some suggest that it has led to the demise of the Fedora Legacy project, which supported older versions of Fedora Core until December of 2006.

The rebuilds of RHEL are freely available; however, you should have a high-speed Internet connection. While these rebuilds do not use 100 percent RHEL software, I have not seen any difference that would impair your ability to study for the Red Hat exams.

- **Community Enterprise Linux** The Community Enterprise Operating System (CentOS) rebuild developed by the group at www.centos.org appears solid to me. This group probably has the largest community (or at least gets the most publicity) among the rebuilds.
- **Scientific Linux** Formerly known as Fermi Linux, it includes a lot of intellectual firepower associated with the Fermi National Accelerator Lab as well as CERN, the lab associated with Tim Berners-Lee, the person most commonly credited with the invention of the World Wide Web.
- **Lineox** Lineox is based in Finland and offers priority updates for a fee. It may be especially interesting for people in the European Union, as their prices are in Euros. You can find out more about Lineox at www.lineox.net.

Alternatively, you can work from RHEL Desktop, if you're willing to install additional services from the source code. For more information on installing packages from source code, see Chapters 1, 5, and 8. Using the techniques described in Chapter 5, you can download the Red Hat Enterprise Linux Source RPMs at ftp.redhat.com, process them into binary RPMs, and then install them on your computer.

For the RHCE exams based on Red Hat Enterprise Linux 5, you can *probably* also work from Fedora Core 6, as RHEL 5 is based on this Red Hat community distribution.

In This Book

The Red Hat RHCT and RHCE exams are designed to test candidate qualifications as Linux systems technicians and engineers. If you pass either of these exams, it's not because you've memorized a canned set of answers-it's because you have a set of Linux administrative skills and know how to use them under pressure, whether it be during an exam or in a real-world situation.

While this book is organized to serve as an in-depth review for the RHCT and RHCE exams for both experienced Linux and Unix professionals, it is not intended as a substitute for Red Hat courses, or more importantly, real-world experience. Nevertheless, each chapter covers a major aspect of the exam, with an emphasis on the "why" as well as the "how to" of working with and supporting RHEL as a systems administrator or engineer. As the actual RHCT and RHCE Exam Prep guide (www.redhat.com/rhce/examprep.html) changes with every release of RHEL (and even sometimes between releases), refer to the noted URL for the latest information. (Throughout the book, I often refer to the RHCT and RHCE Exam Prep guide as the Red Hat Exam Prep guide, even though there are Red Hat exams for certifications other than the RHCT and RHCE.)

Red Hat says it's important to have real-world experience to pass their exams, and they're right! However, for the RHCT and RHCE exams, they do focus on a specific set of Linux administrative skills, as depicted in the Red Hat Exam Prep guide. This book is intended to help you take advantage of the skills you already have-and more importantly, brush up in those areas where you may have a bit less experience.

This book includes relevant information from Red Hat Enterprise Linux 5 (RHEL 5). There are significant changes from Red Hat Enterprise Linux 4; As of this writing, Red Hat even offers a course detailing the differences (RHUP 304 and RHUP 305). Several key differences between RHEL 4 and RHEL 5 include:

- A new hardware detection model. The udev system readily supports automatic mounting and configuration of a wide variety of devices.
- Multicore support. Fundamental to effective virtualization, multicore CPUs can help multiple operating systems run simultaneously on the same physical system. Red Hat includes Xen in RHEL 5 to take advantage of the latest multicore CPUs.
- Logical Volume Management (LVM), version 2, which supports smoother resizing of filesystems.
- Software RAID now supports more modes, including RAID 6. The associated tool is more flexible.
- NFS supports "stateless" network and loopback images.
- yum replaces Up2Date for repository and package management as well as updates.
- The Network Manager incorporates improvements in wireless networking and more, which eases administration on the desktop.
-

SELinux is now easier to use and administer. The descriptions in the Security Level Configuration tool are improved, and **sealert -b** browser can help you diagnose many SELinux-related issues.

There are many more key features; those that I believe are relevant to the RHCT and RHCE exams, as defined by the publicly available course outlines and the Exam Prep guide, are also included in this book.

While it's a risky practice in service, it is fastest to administer RHEL during the exam by logging into the root user account. The command prompt and PATH assume use of that account. When you're logged into the root account, you'll see a command line prompt similar to:

```
[
[root@Enterprise root]#
```

Some Pointers

Once you've finished reading this book, set aside some time to do a thorough review. You might want to return to the book several times and make use of all the methods it offers for reviewing the material:

- *Reread all the Exam Watch notes.* Remember that these notes are written by authors who have taken the exam and passed. They know what you should expect-and what you should be on the lookout for.
- *Review all the Scenario & Solution sections* for quick problem solving.
- *Retake the Self Tests.* Focus on the labs, as there are no multiple choice (or even "fill in the blank") questions on the Red Hat exams. I've included "fill in the blank" questions just to test your mastery of the practical material in each chapter.
- *Complete the exercises.* Did you do the exercises when you read through each chapter? If not, do them! These exercises are designed to cover exam topics, and there's no better way to get to know this material than by practicing. Be sure you understand why you are performing each step in each exercise. If there is something you are not clear on, reread that section in the chapter.

Introduction

The Red Hat Exam Challenge

This section covers the reasons for pursuing industry-recognized certification, explains the importance of your RHCE or RHCT certification, and prepares you for taking the actual examination. It gives you a few pointers on how to prepare, what to expect, and what to do on exam day.

Leaping Ahead of the Competition!

Red Hat's RHCT and RHCE certification exams are hands-on exams. As such, they are respected throughout the industry as a sign of genuine practical knowledge. If you pass, you will be head and shoulders above the candidate who has passed only a "standard" multiple-choice certification exam.

There are two parts to both RHCT and RHCE exams, as follows. The requirements are detailed in the Exam Readiness Checklist later in this introduction.

- **Section I** Troubleshooting and System Maintenance: (RHCE) 2.5 hours (RHCT) 1.0 hours. As described in the Red Hat Exam Prep guide, RHCE candidates need to meet the requirements for both Red Hat certifications. Both RHCT and RHCE candidates must complete *all five RHCT Troubleshooting and System Maintenance problems within the first hour.* As there are five "optional problems," RHCEs need to complete three of five of these RHCE-level problems for an overall score of 80 on this section.
- **Section II** Installation and Configuration: (RHCE) 3.0 hours (RHCT) 2.0 hours. As described in the Red Hat Exam Prep guide, RHCE candidates need to meet the requirements for both the RHCT and RHCE. RHCT candidates need a grade of 70 or higher on their section. RHCE candidates must get a grade of 70 or higher on *both* the RHCT and RHCE components of the Installation and Configuration section of your exam.

Why a Hands-On Exam?

Most certifications today are based on multiple-choice exams. These types of exams are relatively inexpensive to set up and easy to proctor. Unfortunately, many people without real-world skills are good at taking multiple-choice exams. This results in problems on the job with "certified" engineers, who have an image as "paper tigers" who do not have any real-world skills.

In response, Red Hat wanted to develop a certification program that matters. For the most part, they have succeeded with the RHCT, RHCE, and their other advanced certifications.

Linux administrators frequently debug computers with problems. The challenges in the Troubleshooting and System Maintenance sections are based on real-world problems. As the typical Linux administrator has to work through multiple challenges on a daily basis, the RHCE Troubleshooting and System Maintenance section provides a credible measure of real-world skills.

Linux administrators sometimes have to install Linux on a computer or virtual machine. Depending on the configuration, they may need to install Linux from a central source through a network. Installing Linux is not enough to make it useful. Administrators need to know how to configure Linux: add users, install and configure services, create

firewalls, and more.

Exam Watch

The RHCT and RHCE exams are Red Hat exams. Knowledge of System V or BSD-based Unix is certainly helpful, as well as experience with services like Apache, SMB, NFS, DNS, iptables, and DHCP. But it is important to know how to set up, configure, install, and debug these services under Red Hat Enterprise Linux (or rebuild distributions that use the same source code, such as CentOS, Scientific Linux, or Lineox).

Preparing for the RHCT and RHCE Exams

The RHCT is a complete subset of the RHCE. In other words, if you're studying for the RHCT, use this book, based on the guidelines listed in Table 2, at the end of this introduction. If you're studying for the RHCE, read the whole book. The RHCE exam includes every aspect of the RHCT exam.

Work with Red Hat Enterprise Linux. Install it on a computer (or virtual machine) that you don't need for any other purpose. Configure the services described in this book. Find ways to break network services and make Linux unbootable, study the characteristics of the problem, and find different ways to fix the problem.

As you go through this book, you'll have the opportunity to install RHEL several times. If you have more than one computer, you'll be able to install RHEL over a network. And you should, as network installations are explicitly listed in the Exam Prep guide. Then you can work with the different network services. Test out each service as you configure it, preferably from another computer on your network. Testing your work becomes especially important when you start working with the security features of Linux.

Red Hat Certification Program

Red Hat offers several courses that can help you prepare for the RHCT and RHCE. Most of these courses are four or five days long. In some cases, you can take parts of an individual course on an electronic basis. Table 1 illustrates the available hands-on, instructor-led courses that can also help you prepare for the RHCE or RHCT exams.

Table 1: Red Hat RHCT/RHCE Related Courses

Course	**Description**
RH033	Introduction to Linux: basic pre-system administration skills
RH131	Basic system administration skills for installation and configuration (RH133 without the RHCT exam)
RH133	Basic system administration skills for installation and configuration; includes the RHCT exam
RH202	The RHCT exam
RH253	Basic network and security administration; requires a basic knowledge of LANs/WANs and TCP/IP; when combined with RH133, prepares students for the RHCE exam
RH300	The crash course plus the RHCE exam
RH301	The crash course without the RHCE exam
RH302	The RHCE exam
RHUP304	Designed to update RHCEs certified on RHEL 3 and RHEL 4 to RHEL 5, plus the RHCE exam
RHUP305	The update course without the RHCE exam

These aren't the only Red Hat courses available; there are a number of others related to the Red Hat Certified Architect (RHCA) and Red Hat Certified Security Specialist (RHCSS) certifications. But study this first; the RHCE is a prerequisite for those certifications.

Should You Take an RHCT/RHCE Course?

This book is *not* intended as a substitute for the Red Hat RHCE "crash course" (RH300/301) or the RHCT prep course (RH131/RH133). However, the topics are based on the RHCE Rapid Track Course Outline and related RHCT/RHCE Exam Prep guide, both available at www.redhat.com. By design, these topics may help Linux users qualify as administrators and can also be used as such. Just remember, Red Hat can change these topics and course outlines at any time, so monitor www.redhat.com for the latest updates.

RH300 and RH133 are excellent courses. The Red Hat instructors who teach these courses are highly skilled. If you have the skills, it is the best way to prepare for the RHCT and RHCE exams. If you feel the need for classroom instruction, read this book, and then take the appropriate course.

If you're not sure if you're ready for the course or book, read Chapter 1. It is a rapid overview of the prerequisites for the Red Hat RHCT and RHCE certification courses. If you find the material in Chapter 1 to be overwhelming, consider one of the books noted near the start of the chapter, or one of the other Red Hat courses. However, if you are just less familiar with a few of the topics covered in Chapter 1, you're probably okay. Even experienced Linux administrators aren't familiar with everything. Just use the references noted at the beginning of Chapter 1 to fill in any gaps in your knowledge.

Alternatively, you may already be familiar with the material in this book. You may have the breadth and depth of knowledge required to pass the RHCT or RHCE exams. In that case, use this book as a refresher to help you focus on the skills and techniques you need to pass your exam.

Signing Up for the RHCT/RHCE Course and/or Exam

Red Hat provides convenient Web-based registration systems for the courses and test. To sign up for any of the Red Hat courses or exams, navigate to www.redhat.com, click the link for Training and the RHCE/RHCT Program, and select the desired course or exam. Alternatively, contact Red Hat Enrollment Central at (866) 626-2994.

Final Preparations

The Red Hat exams are grueling. Once you have the skills, the most important thing that you can take to the exam is a clear head. If you're tired or frantic, you may miss the easy solutions that are often available. Get the sleep you need the night before the exam. Eat a good breakfast. Bring snacks with you that can keep your mind in top condition.

Remember, the RHCE exam is five and a half hours long. Even the RHCT exam is three hours long. The time allotted for the RHCE exam is more than twice the length of a world-class marathon.

As I describe in Chapter 1, this is an advanced book. It is not designed for beginners to Unix or Linux. As Red Hat does not cover prerequisite skills in its prep course for the RHCT or RHCE exams, I've only covered the tools associated with these prerequisites briefly-mostly in Chapter 1. If you need more information on these prerequisite skills, Red Hat offers other courses (see www.redhat.com/apps/training/); alternatively, read the reference books I've cited in that chapter.

Inside the Exam

The RHCE exam requires that you master RHCT and RHCE skills, and assumes that you already have the prerequisite skills. I've cited them separately, as is done in the current version of the Red Hat Exam Prep guide. Watch for updates at www.redhat.com/training/rhce and www.redhat.com/training/rhce/examprep.html.

Exam RH302

Table 2: Coverage of Red Hat Exam Prep Guide Requirements

Exam Readiness Checklist						
Official Certification Objective	**Study Guide Coverage**	**Ch #**	**Pg#**	**Prerequisite**	**RHCT**	**RHCE**
Red Hat Exam Prerequisite Skills						
Use standard command line tools (e.g., **ls**, **cp**, **mv**, **rm**, **tail**, and **cat**, etc.) to create, remove, view, and investigate files and directories	Basic File Operations and Manipulation	1	19	*		
Use **grep**, **sed**, and **awk** to process text streams and file	Basic File Operations and Manipulation	1	19	*		
Use a terminal-based text editor, such as vim or nano, to modify text files	Basic Linux Knowledge	1	8	*		
Use input/output redirection	Shells	1	26	*		

Understand basic principles of TCP/IP networking, including IP addresses, netmasks, and gateways for IPv4 and IPv6	Basic TCP/IP Networking	1	38	*		
Use **su** to switch user accounts	System Administration	1	34	*		
Use **passwd** to set passwords	Basic Security	1	30	*		
Use **tar**, **gzip**, and **bzip2**	System Administration	1	34	*		
Configure an e-mail client on Red Hat Enterprise Linux	Other Basic Skills as Defined in the Exam Prep Guide	1	49	*		
Use text and/or graphical browser to access HTTP/HTTPS URLs	Other Basic Skills as Defined in the Exam Prep Guide	1	49	*		
Use lftp to access FTP URLs	Other Basic Skills as Defined in the Exam Prep Guide	1	49	*		
RHCT Troubleshooting and System Maintenance Skills					*	

Boot systems into different run levels for troubleshooting and system maintenance	Troubleshooting Strategies	16	728		*	
Diagnose and correct misconfigured networking	Network Configuration	7	331		*	
Diagnose and correct hostname resolution problems	Understanding DNS: Zones, Domains, and Delegation	11	559		*	
Configure the X Window System and a desktop environment	X Window System (entire chapter)	14	649		*	
Add new partitions, filesystems, and swap to existing systems	Partitioning Hard Disks; Managing Filesystems; Advanced Partitioning: Software RAID; Advanced Partitioning: Logical Volume Management	4, 8	185, 196, 410, 417		*	
Use standard command-line tools to analyze problems and configure system	Entire book	all			*	
RHCE Troubleshooting and System Maintenance Skills						

Use the rescue environment provided by first installation CD	Troubleshooting Strategies	16	728			*
Diagnose and correct boot loader failures arising from boot loader, module, and filesystem errors	The GRUB Bootloader, Managing Filesystems, The Basics of the Kernel, Required RHCE Troubleshooting Skills	3, 4, 8, 16	147, 196, 377, 742			*
Diagnose and correct problems with network services (see the following Installation and Configuration skills for a list of these services)		7, 9, 10, 11, 12, 13, 14, 15	329, 443, 493, 557, 585, 613, 649, 691			*
Add, remove, and resize logical volumes	Configuring Partitions, RAID, and LVM; Advanced Partitioning: Logical Volume Management	2, 8	96, 417			*
Diagnose and correct networking service problems where SELinux contexts are interfering with proper operation		7, 9, 10, 11, 12, 13, 14, 15	329, 443, 493, 557, 585, 613, 649, 691			*

RHCT Installation and Configuration Skills						
Perform network OS installation	Configuring a Network Installation	2	81		*	
Implement a custom partitioning scheme	Configuring Partitions, RAID, and LVM	2	96		*	
Configure printing	The CUPS Printing System	7	341		*	
Configure the scheduling of tasks using cron and at	Automating System Administration: cron and at	7	354		*	
Attach system to a network directory service, such as NIS or LDAP	Network Authentication Configuration: NIS and LDAP	6	313		*	
Configure **autofs**	Filesystem Management and the Automounter	4	200		*	
Add and manage users, groups, and quotas, and File Access Control Lists	User Account Management, The Basic User Environment, Setting Up and Managing Disk Quotas	6	273, 285, 290		*	
Configure filesystem permissions for collaboration	Creating and Maintaining Special Groups	6	301		*	

Install and update packages using **rpm**	The Red Hat Package Manager, More RPM Commands	5	222, 227		*	
Properly update the kernel package	New Kernels, the Easy Way	8	388		*	
Configure the system to update/install packages from remote repositories using yum or pup	Adding and Removing RPM Packages with yum and pirut, Managing Updates with Pup and the Red Hat Network (RHN)	5	238, 234		*	
Modify the system boot loader	New Kernels, the Easy Way; Kernel Sources	8	388, 392		*	
Implement software RAID at install-time and runtime	Configuring Partitions, RAID, and LVM; Advanced Partitioning: Software RAID	2, 8	96, 410		*	
Use /proc/sys and **sysctl** to modify and set kernel runtime parameters	The Basics of the Kernel	8	377		*	
Use scripting to automate system maintenance tasks	Automating System Administration: cron and at	7	329		*	

RHCE Installation and Configuration Skills						
For HTTP/HTTPS, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	The Apache Web Server, Virtual Hosts, Apache Access Configuration	9	444, 466, 456			*
For SMB, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	Samba Services	10	516			*
For NFS, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	Configuring a Network File System (NFS) Server, Client-side NFS	10	494, 509			*

For FTP, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	The File Transfer Protocol and vsFTPd	10	512			*
For Web proxy, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	The Squid Web Cache Proxy	9	476			*
For SMTP, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	Electronic Mail (entire chapter)	12	585			*
For IMAP/IMAPS/POP3, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	Reception with Dovecot	12	589			*

For SSH, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	The Secure Shell Package	13	620			*
For DNS (caching name server, slave name server), install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	Domain Name Service (entire chapter)	11	557			*
For NTP, install, configure SELinux support, configure to start on reboot for basic operation and host- and user-based security	The Network Time Protocol (NTP)	13	634			*
Configure hands-free installation using Kickstart	Using Kickstart to Automate Installation	5	244			*
Implement logical volumes at install-time	Configuring Partitions, RAID, and LVM	2	96			*

Use iptables to implement packet filtering and/or NAT	Firewalls and Packet Filtering using netfilter, Network Address Translation	15	697, 703			*
Use PAM to implement user-level restrictions	Pluggable Authentication Modules	6	305			*

Chapter 1: RHCE Prerequisites

Overview

The Red Hat exams are an advanced challenge. As both the RHCE and RHCT courses specify a number of prerequisite skills, this book assumes that you know some basics about Linux. This chapter covers the prerequisite topics for Red Hat's RH300 course in a minimum of detail, with references to other books and sources for more information. It also covers the related prerequisites as defined in the Red Hat Exam Prep guide. Unlike those in other chapters and other books in this series, the questions in this chapter include a number of "zingers" that go beyond the chapter's content. These questions will help determine whether you have the prerequisite skills necessary to handle the remaining chapters.

If you're serious about the RHCE and RHCT exams, this chapter should be just a review. In fact, for any user serious about Linux, this chapter should be trivial. Linux gurus should recognize that I've "oversimplified" a number of explanations; my intention is to keep this chapter as short as possible. However, it is okay if you do not feel comfortable with a few topics in this chapter. In fact, it's quite natural that many experienced Linux administrators don't use every one of the prerequisite topics in their everyday work. Many candidates are able to fill in the gaps in their knowledge with some self-study and practice.

If you're new to Linux or Unix, this chapter will not be enough for you. It's not possible to provide sufficient detail, at least in a way that can be understood by newcomers to Linux and other Unix-based operating systems. If, after reading this chapter, you find gaps in your knowledge, refer to one of the following guides:

- The *Red Hat Enterprise Linux 5* documentation guides, available online from http://www.redhat.com/docs/manuals/enterprise/.
- *Hacking Linux Exposed, Third Edition: Linux Security Secrets and Solutions*, by Casarik, Hatch, Lee, and Kurtz, gives you a detailed look at how to secure your Linux system and networks in every possible way.
- *Mastering Fedora Core 5*, by Michael Jang, covers the distribution that Red Hat used as one of the testbeds for RHEL 5.

Critical to a Linux administrator is knowledge of one or more text editors to manage the many configuration files on a Linux system. The Linux filesystem hierarchy organizes hardware, drivers, directories, and, of course, files. You need to master a number of basic commands to manage Linux. Printer configuration can be a complex topic. Shell scripts enable you to automate many everyday processes. Security is now a huge issue that Linux can handle better than other operating systems, both locally and on larger networks such as the Internet.

As an administrator, you need a good knowledge of basic system administration commands, TCP/IP configuration requirements, and standard network services. While the RHCE and RHCT exams are by and large not hardware exams, some basic hardware knowledge is a fundamental requirement for any Linux administrator.

This is not a book for beginners to Linux/Unix-type operating systems. Some of what you read in this chapter may be unfamiliar. Use this chapter to create a list of topics that you may need to study further. In some cases, you'll be able to get up to speed with the material in other chapters. But if you have less experience with Linux or another Unix-type operating system, you may want to refer to the aforementioned books.

If you're experienced with other Unix-type operating systems such as Solaris, AIX, or HP-UX, you may need to leave some defaults at the door. When Red Hat developed its Linux distribution, it included a number of things that are not consistent with the standards of Unix (or even other Linux distributions). When I took the RH300 course, some students with these backgrounds had difficulties with the course and the RHCE exam.

In this book, most commands are run as the Linux administrative user, root. Logging in as the root user is normally discouraged unless you're administering a computer. However, since the RHCE and RHCT exams test your administrative skills, it's appropriate to run commands in this book as the root user.

While this chapter is based on the prerequisites described at https://www.redhat.com/training/rhce/courses/rh300_prereq.html, there are several additional prerequisite skills defined in the Red Hat Exam Prep guide at https://www.redhat.com/training/rhce/examprep.html.

Inside the Exam

Prerequisite Skills

For the RHCE and RHCT exams, the skills outlined in this chapter are generally minimum requirements. For example, while you may prefer to use an editor other than vi, you may not have access to the GUI, and therefore need to know how to use a console-based text editor on at least the Troubleshooting and System Maintenance section of the exam. While you're not required to know how to pipe the output of dmesg to the **less** command, this is a useful tool that can help you identify problems.

Remember that there are more ways than one to do most everything in Linux. While it's best if you learn all of these "prerequisite" skills, you don't have to know everything in this chapter. In most cases, it's okay if you use other methods of editing or otherwise configuring your RHEL 5 system. As the Red Hat exams no longer include multiple choice questions, don't worry about memorizing the dozens of switches used for certain commands. Focus on results, not trivia.

Using Other Versions of Red Hat

For those of you with more advanced hardware experience, the Red Hat exams are based on PCs built with Intel 32-bit CPUs. That means you'll be using the Linux kernel and associated software that has been customized for this CPU.

For the purpose of this chapter, you can use Fedora Core 6 or one of the rebuild distributions to test your knowledge of basic commands. In fact, the rebuild distributions are excellent, freely available options, as they use the same source code as Red Hat uses to build RHEL 5. One list of rebuild options is available at http://linuxmafia.com/faq/RedHat/rhel-forks.html.

Certification Objective 1.01-Basic Hardware Knowledge

The architecture of a PC defines the components that it uses as well as the way that they are connected. In other words, the Intel-based architecture describes much more than just the CPU. It includes standards for other hardware such as the hard drive, the network card, the keyboard, the graphics adapter, and more. All software is written for a specific computer architecture, such as the Intel-based 32-bit architecture.

Even when a manufacturer creates a device for the Intel platform, it may not work with Linux. Therefore, it's important to know the basic architecture of an Intel-based computer.

Exam Watch

While it is important to know how Linux interacts with your hardware, the RHCE and RHCT exams are not hardware exams. As of this writing, while the RH133 and RH300 courses do address hardware issues, no hardware components are listed in the Red Hat Exam Prep guide.

Architectures

While different versions of RHEL 5 are available for a variety of architectures, you need to be concerned about only one for the Red Hat exams, the basic Intel 32-bit or i386 architecture. As of this writing, the Red Hat exams are offered only on computers with such CPUs, so you need not worry about special architecture-specific issues such as ELILO bootloaders or lib64 module directories.

Intel Communications Channels

Three basic channels are used to communicate in a basic PC: interrupt request (IRQ) ports, input/output (I/O) addresses, and direct memory address (DMA) channels. An IRQ allows a component such as a keyboard or printer to request service from the CPU. An I/O address is a memory storage location for communication between the CPU and different parts of a computer. A DMA channel is used when a device such as a sound card has an independent processor and can bypass the CPU.

With the plug and play features built into RHEL 5, these channels are generally not a problem but are included because they are on the prerequisite list for the RH300 course.

IRQ Settings

An *IRQ* is a signal that is sent by a peripheral device (such as a network card, graphics adapter, mouse, modem, or serial port) to the CPU to request processing time. Each device you attach to a computer may need its own IRQ port. Normally, each device needs a dedicated IRQ (except for USB and some PCI devices).

If you run out of IRQs, some PCI devices can share IRQs. USB devices can share IRQs. This support is available in most PCs manufactured after the year 2000.

On the Job

If you're having a problem with your USB ports or PCI cards, check your BIOS first. Many BIOS menus include options that enable PCI sharing and support USB connections.

Planning the IRQ Layout: Standard IRQs

IRQs are a precious commodity on some PCs. IRQ conflicts are common when you're connecting a lot of devices. If your printer doesn't work after you've connected a second network card, it can help to know the standard IRQ for printers. You can then assign a different IRQ to that network card. If you don't have any free IRQs to assign to that network card, you may be able to sacrifice a component that uses a standard IRQ. For example, if you always connect to a server remotely, that server PC may not need a keyboard. If you can boot a computer with a CD-ROM, you may not need a floppy drive.

Some IRQs are essential to the operation of a PC and just can't be changed. These are reserved by the motherboard to control devices such as the hard disk controller and the real-time clock. Do not use these interrupts for other devices or there will be conflicts! Other IRQs are normally assigned to common devices such as a floppy drive and a printer. In Linux, you can check /proc/interrupts to see which interrupts are being used and which are free for new devices.

Input/Output Addresses

Every computer device requires an *input/output (I/O) address.* It's a place where data can wait in line for service from your CPU. I/O addresses are listed in hexadecimal notation, where the "numbers" are 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, a, b, c, d, e, and f. Some typical I/O addresses include those for the basic serial ports, known in the Microsoft world as COM1, COM2, COM3, and COM4. These ports normally use the following I/O addresses: 03f8, 02f8, 03e8, and 02e8.

You can find a list of assigned I/O addresses in your /proc/ioports file.

Direct Memory Addresses

A *direct memory address (DMA)* is normally used to transfer information directly between devices, bypassing the CPU. Many components don't need a CPU. For example, many sound cards include their own processor. This allows your PC to set up a DMA channel between a hard drive and a sound card to process and play any music files that you may have stored.

While DMA channels bypass the CPU, devices that use DMA are still configured with IRQ ports. There are eight standard DMA channels (0?7); DMA 4 is reserved and cannot be used by any device.

You can find a list of assigned DMA addresses in your /proc/dma file.

RAM Requirements

While I've installed RHEL 5 on computers with less RAM, 256MB is a good practical minimum.

The maximum amount of memory your system will use is the sum of all of the memory requirements of every program that you will ever run at once. That's hard to compute. Therefore, you should buy as much memory as you can afford. Extra RAM is usually cost effective when compared to the time you would spend trying to tune an underpowered system. Limitations are few; on Red Hat Enterprise Linux 5 Advanced Platform (with PAE support), RAM is limited only by the ability of the hardware to handle it.

On the Job

If you're setting up Linux as a server, RAM requirements increase with the number of users who may need to log in simultaneously. The same may be true if you're running a large number of programs or have memory-intensive data such as that required by a database.

Hard Drive Options

Before your computer can load Linux, the BIOS has to recognize the active primary partition on the hard drive. This partition should include the Linux boot files. The BIOS can then set up and initialize that hard drive, and then load Linux boot files from that active primary partition. You should know the following about hard drives and Linux:

- The standard PC is configured to manage up to four IDE (Integrated Drive Electronics) hard drives, now known as PATA (Parallel Advanced Technology Attachment) drives.
- Newer PCs can handle more SATA (Serial ATA) drives.
- Depending on the SCSI (Small Computer Systems Interface) hardware that you have, you can attach up to 31 different SCSI hard drives.
- While you can use as many PATA, SATA, or SCSI drives as your hardware can handle, you need to install the Linux boot files from the /boot directory on one of the first two hard drives. If Linux is installed on a later drive, you'll need a boot floppy.
- Although you can install Linux on USB (Universal Serial Bus) or IEEE 1394 (Institute of Electrical and Electronics Engineers standard 1394, also known as FireWire or iLink) hard drives, as of this writing, you can't load Linux boot files directly from these drives. However, it is possible to set up a boot floppy or CD/DVD to start Linux from these drives.

Certification Objective 1.02-Basic Linux Knowledge

Linux and Unix are managed through a series of text files. Linux administrators do not normally use graphical editors to manage these configuration files. Editors such as WordPerfect, OpenOffice.org Writer, and yes, even Microsoft Word normally save files in a binary format that Linux can't read.

Popular text editors for Linux configuration files include nano, pico, joe, and vi. If you already know one of these editors, feel free to skip this section. If you have to rescue an RHEL 5 system (as may be required during the exam), you'll have access to these editors when booting your system from RHEL 5 rescue media.

The VIsual Editor

While emacs may be the most popular and flexible text editor in the world of Linux, I believe every administrator needs at least a basic knowledge of vi, which may help you save a broken system. If you ever have to restore a critical configuration file using an emergency boot floppy, vi is probably the only editor that you'll have available.

In reality, RHEL 5 uses an enhanced version of the vi editor, known as vim. And as RHEL emergency boot media access installation packages, it supports more console-based editors. I describe vi here simply because it's the editor I know best.

On the Job

If you boot in rescue mode and try to start emacs or pico, that starts the joe editor instead.

You should know how to use the two basic modes of vi: command and insert. When you use vi to open a file, it opens in command mode. Some of the commands start insert mode. Opening a file is easy: just use the vi ***filename*** command. By default, this starts vi in command mode. An example of vi with the /etc/inittab file is shown in Figure 1-1.

```
#
# inittab       This file describes how the INIT process should set up
#               the system in a certain run-level.
#
# Author:       Miquel van Smoorenburg, <miquels@drinkel.nl.mugnet.org>
#               Modified for RHS Linux by Marc Ewing and Donnie Barnes
#

# Default runlevel. The runlevels used by RHS are:
#   0 - halt (Do NOT set initdefault to this)
#   1 - Single user mode
#   2 - Multiuser, without NFS (The same as 3, if you do not have networking)
#   3 - Full multiuser mode
#   4 - unused
#   5 - X11
#   6 - reboot (Do NOT set initdefault to this)
#
id:5:initdefault:

# System initialization.
si::sysinit:/etc/rc.d/rc.sysinit

l0:0:wait:/etc/rc.d/rc 0
"/etc/inittab" [readonly] 53L, 1663C
```

Figure 1-1: The vi editor with/etc/inittab

The following is only the briefest of introductions to the vi editor. For more information, there are a number of books available, as well as an extensive manual formatted as a HOWTO available from the Linux Documentation Project at www.tldp.org. Alternatively, a tutorial is available through the **vimtutor** command.

vi Command Mode

In command mode, you can do everything you need to a text file except edit it. The options in command mode are

broad and varied, and they are the subject of a number of book-length texts. In summary, vi requires seven critical command skills:

- **Open** To open a file in the vi editor from the command line interface, run the vi ***filename*** command.
- **Search** Start with a backslash, followed by the search term. Remember, Linux is case-sensitive, so if you're searching for "Michael" in /etc/passwd, use the **/Michael** (not **/michael**) command.
- **Write** To save your changes, use the **w** command. You can combine commands: for example, **:wq** writes the file and exits vi.
- **Close** To leave vi, use the **:q** command.
- **Abandon** If you want to abandon any changes that you've made, use the **:q!** command.
- **Edit** You can use a number of commands to edit files through vi, such as **x**, which deletes the currently highlighted character; **dw**, which deletes the currently highlighted word; and **dd**, which deletes the current line. Remember, **p** places text from a buffer, and **U** restores text from a previous change.
- **Insert** A number of commands allow you to start insert mode, including **i** to start inserting text at the current position of the editor, and **o** to open up a new line immediately below the current position of the cursor.

Basic Text Editing

In modern Linux systems, editing files with vi is easy. Just use the normal navigation keys (arrow keys, PAGE UP, and PAGE DOWN), and then one of the basic commands such as **i** or **o** to start vi's insert mode, and type your changes directly into the file.

When you're finished with insert mode, press the ESC key to return to command mode. You can then save your changes or abandon them and exit vi.

On the Job

There are several specialized variations on the vi command. Three are *vipw*, *vigw*, and *visudo*, which edit /etc/passwd, /etc/group, and /etc/sudoers, respectively.

Exercise 1-1: Using vi to Create a New User

In this exercise, you'll create a new user by editing the /etc/passwd file with the vi text editor. While you could create new Linux users in other ways, this exercise helps you verify your skills with vi and at the command line interface.

1. Open a Linux command line interface. Log in as the root user, and type the **vipw** command. This command uses the vi editor to open /etc/passwd.
2. Navigate to the last line in the file. As you should already know, there are several ways to navigate in command mode, including the DOWN ARROW key, the PAGE DOWN key, the **G** command, or even the

K key.

3. Make one copy of this line. If you're already comfortable with vi, you should know that you can copy an entire line to the buffer with the **yy** command. This "yanks" the line into buffer. You can then restore or "put" that line as many times as desired with the **p** command.

4. Change the username, user ID, group ID, user comment, and home directory for the new user. If you understand the basics of Linux or Unix, you'll understand their locations on each line in the /etc/passwd file. For example, in Figure 1-2, this corresponds to gb, 501, 501, Gordon Brown, and /home/gb. Make sure the username also corresponds to the home directory.

5. Return to command mode by pressing the ESC key. Save the file with the **:w** command, and then exit with the **:q** command. (You can combine the two commands in vi; the next time you make a change and want to save and exit, run the **:wq** command.)

6. As the root user, run the **passwd** ***newuser*** command. Assign the password of your choice to the new user.

```
netdump:x:34:34:Network Crash Dump user:/var/crash:/bin/bash
avahi:x:70:70:Avahi daemon:/:/sbin/nologin
named:x:25:25:Named:/var/named:/sbin/nologin
mailnull:x:47:47::/var/spool/mqueue:/sbin/nologin
smmsp:x:51:51::/var/spool/mqueue:/sbin/nologin
haldaemon:x:68:68:HAL daemon:/:/sbin/nologin
rpc:x:32:32:Portmapper RPC user:/:/sbin/nologin
rpcuser:x:29:29:RPC Service User:/var/lib/nfs:/sbin/nologin
nfsnobody:x:4294967294:4294967294:Anonymous NFS User:/var/lib/nfs:/sbin/nologin
sshd:x:74:74:Privilege-separated SSH:/var/empty/sshd:/sbin/nologin
xfs:x:43:43:X Font Server:/etc/X11/fs:/sbin/nologin
beagleindex:x:58:58:User for Beagle indexing:/var/cache/beagle:/bin/false
distcache:x:94:94:Distcache:/:/sbin/nologin
ntp:x:38:38::/etc/ntp:/sbin/nologin
squid:x:23:23::/var/spool/squid:/sbin/nologin
dovecot:x:97:97:dovecot:/usr/libexec/dovecot:/sbin/nologin
mysql:x:27:27:MySQL Server:/var/lib/mysql:/bin/bash
webalizer:x:67:67:Webalizer:/var/www/usage:/sbin/nologin
hsqldb:x:96:96::/var/lib/hsqldb:/sbin/nologin
gdm:x:42:42::/var/gdm:/sbin/nologin
postfix:x:89:89::/var/spool/postfix:/sbin/nologin
michael:x:500:500::/home/michael:/bin/bash
gb:x:501:501:Gordon Brown:/home/gb:/bin/bash
                                                  23,1          Bot
```

Figure 1-2: Adding a new user in /etc/passwd

Other Text Editors

You can use any available text editor during the Red Hat exams. But do not count on any GUI text editors, as the GUI may not be available for troubleshooting during the exam, or in real life.

Certification Objective 1.03-Linux Filesystem Hierarchy and Structure

Everything in Linux can be reduced to a file. Partitions are associated with filesystem *device nodes* such as /dev/hda1. Hardware components are associated with node files such as /dev/dvd. Detected devices are documented as files in the /proc directory. The Filesystem Hierarchy Standard (FHS) is the official way to organize files in Unix and Linux directories. As with the other sections, this introduction provides only the most basic overview of the FHS. More information is available from the official FHS homepage at www.pathname.com/fhs.

Linux Filesystems and Directories

Several major directories are associated with all modern Unix/Linux operating systems. These directories organize user files, drivers, kernels, logs, programs, utilities, and more into different categories. The standardization of the FHS makes it easier for users of other Unix-based operating systems to understand the basics of Linux.

Every FHS starts with the root directory, also known by its symbol, the single forward slash (/). All of the other directories shown in Table 1-1 are subdirectories of the root directory. Unless they are mounted separately, you can also find their files on the same partition as the root directory. You may not see some of the directories shown in the table if you have not installed associated packages. Not all directories shown are officially part of the FHS.

Table 1-1: Basic Filesystem Hierarchy Standard Directories

Directory	**Description**
/	The root directory, the top-level directory in the FHS. All other directories are subdirectories of root, which is always mounted on some partition.
/bin	Essential command line utilities. Should not be mounted separately; otherwise, it could be difficult to get to these utilities when using a rescue disk.
/boot	Includes Linux startup files, including the Linux kernel. The default, 100MB, is usually sufficient for a typical modular kernel and additional kernels that you might install during the RHCE or RHCT exam.
/dev	Hardware and software device drivers for everything from floppy drives to terminals. Do not mount this directory on a separate partition.
/etc	Most basic configuration files.
/home	Home directories for almost every user.
/lib	Program libraries for the kernel and various command line utilities. Do not mount this directory on a separate partition.

/media	The mount point for removable media, including floppy drives, DVDs, and Zip disks.
/misc	The standard mount point for local directories mounted via the automounter.
/mnt	A legacy mount point; formerly used for removable media.
/net	The standard mount point for network directories mounted via the automounter.
/opt	Common location for third-party application files.
/proc	Currently running kernel-related processes, including device assignments such as IRQ ports, I/O addresses, and DMA channels, as well as kernel configuration settings such as IP forwarding.
/root	The home directory of the root user.
/sbin	System administration commands. Don't mount this directory separately.
/selinux	Currently configured settings associated with Security Enhanced Linux.
/smb	The standard mount point for remote shared Microsoft network directories mounted via the automounter.
/srv	Commonly used by various network servers on non-Red Hat distributions.
/tftpboot	Included if the TFTP server is installed.
/tmp	Temporary files. By default, Red Hat Enterprise Linux deletes all files in this directory periodically.
/usr	Small programs accessible to all users. Includes many system administration commands and utilities.
/var	Variable data, including log files and printer spools.

Mounted directories are often known as *volumes*, which can span multiple partitions. However, while the root directory (/) is the top-level directory in the FHS, the root user's home directory (/root) is just a subdirectory.

On the Job

In Linux, the word filesystem has several different meanings. For example, a filesystem can refer to the FHS, an individual partition, or a format such as ext3. A filesystem device node such as /dev/sdal represents the partition on which you can mount a directory.

A Variety of Media Devices

Several basic types of media are accessible to most PCs, including PATA, SATA, and SCSI hard disks; floppy drives; DVD/CD drives; and more. Other media are accessible through other PC ports, including serial, parallel, USB, and IEEE 1394 systems. You can use Linux to manage all of these types of media.

Most media devices are detected automatically. Linux may require a bit of help for some devices described in Chapter 2. But in the context of the Linux FHS, media devices, like all others, are part of the /dev directory. Typical media devices are described in Table 1-2.

Table 1-2: Media Devices

Media Device	**Device File**
Floppy drive	First floppy (Microsoft A: drive) = /dev/fd0 Second floppy (Microsoft B: drive) = /dev/fd1
PATA (IDE) hard drive PATA (IDE) CD/DVD drive	First drive = /dev/hda Second drive = /dev/hdb Third drive = /dev/hdc Fourth drive = /dev/hdd
SATA or SCSI hard drive SATA or SCSI CD/DVD drive	First drive = /dev/sda Second drive = /dev/sdb ? Twenty-seventh drive = /dev/sdaa and so on
Parallel port drives	First drive = /dev/pd1 First tape drive: /dev/pt1
USB drives	Varies widely
IEEE 1394 drives	IEEE 1394 (a.k.a. FireWire, iLink) is actually a SCSI standard, so these are controlled in Linux as SCSI devices.

Making Reference to Devices in /dev

Take a look at the files in the /dev directory. Use the **ls -l /dev | more** command. Scroll through the list for a while. The list actually used to be longer. Well, there's a method to this madness. Some devices are linked to others, and that actually makes it easier to understand what is connected to what. For example, the virtual device files /dev/cdrom and /dev/dvd are easier to identify than the true device files. Generally, these devices are automatically linked to the actual device files during Linux installation. For example, if you have a printer and DVD writer installed, the following commands illustrate possible links between these components and the actual device files:

```
#
# ls -l /dev/par0
 l
 lrwxrwxrwx 1 root root 3 Mar 29 09:37 /dev/par0 -> lp0
#
```

Certification Objective 1.04-Basic File Operations and Manipulation

Linux was developed as a clone of Unix, which means that Linux has the same functionality with different source code. The essence of both operating systems is at the command line. Basic commands for file manipulation and filters are available to help you do more with a file.

Exam Watch

This section covers only the most basic of commands that you can use in Linux. It describes only a few of the things that you can do with each command. Unfortunately, a full discussion would require several hundred more pages. Expect to know considerably more about commands for the RHCE and RHCT exams.

Basic File Operations

Two basic groups of commands are used to manage Linux files. One group helps you get around Linux files and directories. The other group actually does something creative with the files. Remember that in any Linux file operation, you can take advantage of the **HISTORY** (this is capitalized because it's a standard environment variable) of previous commands, as well as the characteristics of command completion, which allow you to use the TAB key almost as a wildcard to complete a command or a filename or give you the options available in terms of the absolute path.

Almost all Linux commands include *switches*, options that allow you to do more. Few are covered in this chapter. If you're less familiar with any of these commands, use their man pages. Study the switches. Try them out! Only with practice, practice, and more practice can you really understand the power behind some of these commands.

Basic Navigation

Everything in Linux can be reduced to a file. Directories are special types of files that serve as containers for other files. Drivers are files. As discussed earlier, devices are special types of files. The nodes associated with USB hardware are just files, and so on. To navigate around these files, you need some basic commands to tell you where you are, what is there with you, and how to move around.

The Tilde (~)

But first, every Linux user has a home directory. You can use the tilde (~) to represent the home directory of any currently active user. For example, if your username is tb, your home directory is /home/tb. If you've logged in as the root user, your home directory is /root. Thus, the effect of the **cd ~** command depends on your username. For example, if you've logged in as user mj, the **cd ~** command brings you to the /home/mj directory. If you've logged in as the root user, this command brings you to the /root directory. You can list the contents of your home directory from anywhere in the directory tree with the **ls ~** command.

Paths

There are two path concepts associated with Linux directories: absolute paths and relative paths. An absolute path describes the complete directory structure based on the top level directory, root (/). A relative path is based on the current directory.

Relative paths do not include the slash in front.

The difference between an absolute path and a relative one is important. Especially when you're creating a script, absolute paths are essential. Otherwise, scripts executed from other directories may lead to unintended consequences.

pwd

In many configurations, you may not know where you are relative to the root (/) directory. The **pwd** command, which is short for print working directory, can tell you, relative to root (/). Once you know where you are, you can determine whether you need to move to a different directory.

cd

It's easy to change directories in Linux. Just use **cd** and cite the absolute path of the desired directory. If you use the relative path, just remember that your final destination depends on the present working directory.

ls

The most basic of commands lists the files in the current directory. But the Linux **ls** command, with the right switches, can be quite powerful. The right kind of **ls** can tell you everything about a file, such as creation date, last access date, and size. It can help you organize the listing of files in just about any desired order. Important variations on this command include **ls -a** to reveal hidden files, **ls -l** for long listings, **ls -t** for a time-based list, and **ls -i** for inode numbers. You can combine switches; I often use the **ls -ltr** command to display the most recently changed files last.

Looking for Files

There are two basic commands used for file searches: find and locate.

find

The find command searches through directories and subdirectories for a desired file. For example, if you wanted to find the directory with the xorg.conf GUI configuration file, you could use the following command, which would start the search in the top-level root (/) directory:

```
#
# find / -name xorg.conf
```

Certification Objective 1.05-Printing

As of this writing, printers are not always connected or configured during the installation of Red Hat Enterprise Linux. You may have to install printers yourself. The default Red Hat Enterprise Linux print daemon is CUPS, the Common Unix Printing System.

There are three basic ways to configure a printer: first, you can edit the configuration files in the /etc/cups directory with a text editor, which can be a difficult process. These files are long, and the language is somewhat obscure, at least on the surface.

Through its support of the Internet Printing Protocol (IPP), CUPS provides another way toward managing printers on a network: a Web-based configuration tool using TCP/IP port 631.

The third method in RHEL 5 is with the Red Hat Printer Configuration tool, which is described in Chapter 7.

Adding Printers

The easy way to add a printer is with the Red Hat Printer Configuration tool, which is also known by the command used to start it from a terminal, **system-config-printer**. I recommend that you learn to use this GUI tool. Unless you're a CUPS expert, it's a faster way to configure printers on the RHCT and RHCE exams. I show you how to use this utility in Chapter 7.

Print Commands

Three basic commands are associated with printing in Linux, as described in Table 1-5.

Table 1-5: Linux Print Commands

Command	**Description**
lpr	The basic print command; lpr ***filename*** prints that file.
lpq	Query the print queue for status; **lpr -l** lists print job numbers.
lprm	Remove a specific job, usually specified by job number, from the printer queue.

Certification Objective 1.06-Shells

A *shell* is a user interface. The Linux command shell is the prompt that allows you to interact with your computer with various system commands. With the right file permissions, you can set up commands in scripts to run when you want, even in the middle of the night. Linux shells can process commands in various sequences, depending on how you manage the input and output of each command. The way commands are interpreted is in part determined by variables and parameters associated with each shell. Some of these variables make up the environment that is carried over even if you change from one shell to another.

The default shell in Linux is bash, also known as the Bourne Again Shell. A number of other shells are popular with many users. As long as you have installed the appropriate RPMs, users can start any of these shells. As desired, you can change the default shell for individual users in the /etc/passwd file.

Basic Shell Programming

"Real" Linux administrators program their own scripts. They create scripts because they don't want to sit at their computers all the time. Scripts can allow Linux to back up directories automatically when nobody is in the office. Scripts can help Linux process databases when few people are using the system.

If you're not a programmer, don't worry-this is not as difficult as it sounds. For example, utilities related to the crontab command automate the creation of a number of different scripts. The cron system is discussed in more detail in Chapter 6.

If you're at all familiar with shell commands and programming expressions, you can find some examples of Red Hat Enterprise Linux shell programs in the /etc/cron .daily directory.

Variables and Parameters

Variables can change. Parameters are set. The bash shell includes a number of standard environment variables. Their default values are shown in the output to the **env** command. One critical variable is the value of PATH, which you can check at the command line with the **echo $PATH** command. The directories listed in PATH are automatically searched when you try to run a command. For example, if you want to run the fdisk command from the /sbin directory, you could do it with the following command:

```
$
$ /sbin/fdisk
```

Certification Objective 1.07-Basic Security

The basic security of a Linux computer is based on file permissions. Default file permissions are set through the umask shell variable. SUID and SGID permissions can give all users access to specific files. Ownership is based on the default user and group IDs of the person who created a file. Managing permissions and ownership involves commands such as chmod, chown, and chgrp.

Users and groups own files. Users and groups have passwords. Security can be enhanced if you configure users and groups in the Shadow Password Suite. Obviously, more levels of security are available, but security options such as Access Control Lists and Security Enhanced Linux (SELinux) are not included in the Red Hat exam prerequisites.

File Permissions

Linux file permissions are straightforward. Consider the following output from **ls -l /sbin/fdisk**:

```
-rwxr-xr-x 1 root root 95572 Jan 11 08:10 /sbin/fdisk
```

Certification Objective 1.08-System Administration

Most system administration tasks require root or superuser privileges. You should already be familiar with a number of basic Linux system administration commands and files. Standard user files are stored in /etc/skel. Daemons are processes that run in the background and run various Linux services. cron is a specialized daemon that can run scripts when you want. It's especially useful for setting up backup jobs in the middle of the night. Logging is a key part of monitoring Linux and any services that you choose to run.

The Superuser

Generally in Linux, a system administrator does everything possible as a normal user. It's a good practice to use superuser privileges only when absolutely necessary. But one time when it's appropriate is during the Red Hat exams. Good administrators will return to being normal users when they're done with their tasks. Mistakes as the root user can disable your Linux system.

There are two basic ways to make this work:

- **su** The superuser command, **su**, prompts you for the root password before logging you in with root privileges. A variation, **su -c**, sets up root privileges for one specific command. Many Red Hat GUI utilities are set up to prompt for the root password before they can be started using Pluggable Authentication Modules (see Chapter 6). One more variation, **su - root**, sets up root privileges with the root user PATH. (Remember to use a space on both sides of the dash in this command.)
- **sudo** The **sudo** command allows users listed in /etc/sudoers to run administrative commands. You can configure /etc/sudoers to set limits on the root privileges granted to a specific user.

However, Red Hat Enterprise Linux provides some features that make working as root somewhat safer. For example, logins using the **ftp** and telnet commands to remote computers are disabled by default.

/etc/skel for Home Directories

Basic configuration files for individual users are available in the /etc/skel directory. This directory includes a number of hidden files. For a full list, run the **ls -a /etc/skel** command. If you want all future users to get specific files in their home directories, include them here.

The next time you create a regular user, check that person's home directory. For example, if you just created a user named elizabeth, run the **ls -a /home/elizabeth** command. Compare the results to the previous command on the /etc/skel directory.

Daemons

A daemon is a process that runs in the background. It is resident in your computer's RAM and watches for signals before it goes into action. For example, a network daemon such as httpd, the Linux Web server known as Apache, waits for a request from a browser before it actually serves a Web page.

Daemons are often configured to start automatically when you start Linux. This process is documented at various runlevels in the /etc/rc.d directory. Alternatively, you can use a tool such as ntsysv to identify and manage the daemons that are started at various Linux runlevels. This is discussed in more detail in Chapter 4.

Controlling Network Services Through Daemons

Networks don't always work. Sometimes you need to restart a network daemon to implement a configuration change. Red Hat Enterprise Linux provides an easy way to control network service daemons through the scripts in /etc/rc.d/init.d. This directory includes scripts that can control installed Linux network services (and more) for everything from the Network File System (NFS) to sendmail. The actual daemon itself is usually located in the /sbin or /usr/sbin directory.

Exam Watch

In Red Hat Enterprise Linux, a simpler way to reload or restart a service in the /etc/init.d directory is with the service command. For example, to restart the vsftpd service, you could run the service vsftpd restart command. (And that's one more reason to log in as the root user; if you invoke root privileges with su, based on the default $PATH, you'd have to type /sbin/service vsftpd restart.)

With these scripts, it's easy to start, stop, status, reload, or restart a network daemon. This is useful to implement or test changes that you make to a specific configuration file. For example, if you make a change to the Postfix mail server configuration file in /etc/ postfix/main.cf, you can implement the change right away with the **/etc/init.d/postfix reload** command. Other switches to these scripts allow you to stop, start, or status these services. Service management is discussed in more detail in Chapter 3.

cron

Perhaps the most important daemon is cron, which can be used to execute a command or a series of commands in a script, on a schedule. Red Hat Enterprise Linux already includes a series of scripts that are executed by cron on committed schedules in the /etc/cron.hourly, /etc/cron.daily, /etc/cron.weekly, and /etc/cron .monthly directories.

System crontab

The easiest way to set up your own cron jobs is through the crontab file, which can be managed through the crontab command. Users can edit their own crontab files with the **crontab -e** command; the root user can configure the crontab for a specific user with the **crontab -u *username* -e** command.

The general format for a crontab file can be found in the /etc/crontab script, which is used to run the scripts in the aforementioned schedule-related directories. A typical crontab entry from that file is

```
4
42 4 1 * * root run-parts /etc/cron.monthly
```

Certification Objective 1.09-Basic TCP/IP Networking

TCP/IP is a series of protocols organized in layers, known as a protocol suite. It was developed for Unix and eventually adopted as the standard for communication on the Internet. With IP addresses, it can help you organize your network. There are a number of TCP/IP tools and configurations that can help you manage your network.

As with the previous sections in this chapter, the statements here are oversimplifications. So if you find this section overwhelming, read the references cited at the beginning of the chapter. Linux is built for networking, and there is no practical way to pass either the RHCT or the RHCE exam unless you understand networking in some detail.

IP Numbers and Address Classes

Every computer that communicates on a network needs its own IP address. Some addresses are assigned permanently to a particular computer; these are known as *static* addresses. Others are leased from a DHCP server, associated with the Dynamic Host Configuration Protocol, for a limited amount of time; these are also known as *dynamic* IP addresses.

Two standards for IP addresses are in use today: IP version 4 (IPv4) and IP version 6 (IPv6). IPv4 addresses have 32 bits and are set up in octets in dotted decimal notation. The range of possible IPv4 addresses is between 0.0.0.0 to 255.255.255.255. While this range includes more than 4 billion IP addresses, that is not nearly enough for the current Internet.

IPv6 addresses have 128 bits and are set up in hexadecimal notation. An IPv6 address is normally organized in eight groups of four hexadecimal numbers each, and it may look like *4abe:03e2:c132:69fa:0000:0000:c0b8:2148*. This is a range of over 340,000,000,000,000,000,000,000,000,000,000,000,000 IPv6 addresses.

To ease the transition, specific IPv6 addresses have been assigned for every one of the 4 billion IPv4 addresses. There are still more than 3.4 × 1038 addresses left over. While actual routing on the Internet now commonly uses IPv6, network configuration in Linux is still normally based on IPv4 addresses.

IPv4 addresses are organized into five different classes, as shown in Table 1-7. The academics among you may note that this is different from the official addresses in each IPv4 class as specified in RFC 1518 from the Internet Engineering Task Force (www.ietf.org). The *assignable* address range includes those IP addresses that can be assigned to a specific computer on a network.

Table 1-7: IP Address Classes

Class	Assignable Address Range	Note
A	1.1.1.1?126.255.255.254	Allows networks of up to 16 million computers
B	128.0.0.1?191.255.255.254	Allows networks of up to 65,000 computers
C	192.0.0.1?223.255.255.254	Allows networks of up to 254 computers

D	224.0.0.1?239.255.255.254	Reserved for multicasts
E	240.0.0.1?255.255.255.254	Reserved for experimental use

In addition, there are several private IP address ranges available for computers and networks. They are associated with network addresses 10.0.0.0, 172.168.0.0, and 192.168.0.0 through 192.168.255.0.

IPv6 Addressing

There are 128 bits in an IPv6 address. That's 96 more bits than IPv4. A typical IPv6 address might look like this:

```
6
63 .51 .33 .1 67 .11 0.5 0. 0. 0. 0. 0.153 .53 .68.9 8.31 . 235
```

Certification Objective 1.10-Familiarity with Standard Network Services

Linux is built for networking. The code associated with many standard networking services is integrated into the Linux kernel. A basic understanding of the functionality of standard Linux networking services is essential. Many themes throughout this book assume that you already understand the purposes of network communication protocols, mail services, host name and IP address management, Web services, and more.

In Red Hat Enterprise Linux, network services are often installed separately. Some include different packages for clients and servers. Some network services are activated through /etc/xinetd.conf, which reads activation files in the /etc/xinetd.d directory. Others are activated directly with scripts in the /etc/init.d directory. Some key RHEL network services are briefly examined in the following sections.

Network File System, Locally and Remotely

The first network system on Unix and Linux computers is NFS. Ideally, this leads to a seamless Linux interface; for example, you can set up one /home directory for all users on your network on one server. Remember that you need NFS on both server and client computers on your network.

First, make sure NFS support is part of the kernel, as documented in /proc/ filesystems. If it isn't there, you may need to activate the nfs and related modules (nfsd, lockd, sunprc) in the kernel. Inspect installed modules with the following command:

```
#
# lsmod | more
```

Certification Objective 1.11-Basic Network Security

Network security in Linux has five basic components. Security by computer can help you manage what computers can send messages into and out of your network. Security by port can help you manage the services that others can use to break into your network. Security by address translation can help you hide the computers inside your network. Security by rule can help you manage the type of data allowed into your network in excruciating detail. And finally, security by SELinux can help manage network services on an entirely different level. Red Hat Enterprise Linux includes tools that can help you configure a firewall and SELinux on your computer: **system- config-securitylevel** (also known as the Red Hat Security Level Configuration tool) and **system-config-selinux** (also known as the SELinux Management Tool). Firewalls are covered in Chapter 15 and SELinux configuration is covered in Chapters 4 and 15. As SELinux is not a "prerequisite" skill, it is not covered in this chapter.

Allowing and Denying

The /etc/hosts.allow and /etc/hosts.deny files can help you manage what computers are allowed into your network. You can specify computers by name, IP address, network, or domain name in each file. This can help you limit access to a trusted few computers such as those within your company, or it can protect you from computers that you know may pose a problem.

Securing Ports

TCP/IP has 65,536 ports, which work sort of like TV channels. If you leave all ports open, you're leaving a lot of options for a cracker who wants to break into your network. With a firewall, you can create a solid barrier and then open only the ports that you need.

Network Address Translation

Most LAN administrators set up Network Address Translation (NAT) as a matter of course on an IPv4 network. Since IPv4 addresses are scarce, it is typical to use private IP addresses inside a LAN, with a regular IP address only on the gateway computer that is directly connected to an outside network such as the Internet.

For example, when a computer inside a LAN wants access to a Web page, NAT sends the IP address of the gateway to the Internet. Nobody outside the LAN need know the real source of the Web page request.

iptables

The iptables command has three basic ways to look at a data packet: input, output, or forward. Within these and other parameters, you can set up your firewall with instructions to let the packet pass, let it drop, or direct it someplace else. If you're working with IPv6, the corresponding command is **ip6tables**.

Once you've configured a firewall and loaded it, the rules are stored in the /etc/ sysconfig/iptables file. iptables is covered in more detail in Chapter 15.

Certification Objective 1.12-Other Basic Prerequisite Skills per the Red Hat Exam Prep Guide

There are more prerequisite skills defined in the Red Hat Exam Prep guide at https://www.redhat.com/training/rhce/examprep.html. These are over and above the RH300 prerequisites as defined by https://www.redhat.com/training/rhce/courses/rh300_prereq.html. Word for word from the current Exam Prep guide, they specify that you know how to:

- Configure an e-mail client on Red Hat Enterprise Linux
- Use a text and/or graphical browser to access HTTP/HTTPS URLs
- Use lftp to access FTP URLs

Configuring an email Client

The configuration process for a GUI e-mail client should be trivial for any candidate for Red Hat certification. However, the same may not necessarily be true for command line clients, and it's certainly possible that you'll have to configure RHEL 5 solely on the command line.

Command Line Mail

To test your mail system, you can use the built-in command line **mail** utility, a simple text-based interface. The system keeps each user's mail in a system directory. Once users read a message, they can reply, forward, or delete it. If they do not delete the message before quitting the mail utility, the system stores the message in the /var/ mail directory, in a file named after the applicable username.

You can certainly use any of the other mail readers, such as mutt, or the e-mail managers associated with different GUI Web browsers to test your system. Other mail readers store messages in different directories. For example, pine would create and store messages for user mj in the /home/mj/mail directory.

To send mail to another user, you can use the mail command line utility. There are two basic methods for using mail. First, you can enter the subject and then the text of your message. When you're done, press CTRL-D and then enter another addressee in the Cc: line, if desired. When you press ENTER, the message is sent and the mail utility stops and sends you back to the command line:

```
 $
 $ mail Michael
S
Subject: Test Message
```

Certification Objective 1.13-Downloading the Red Hat Enterprise Linux Installation CDs

First, this section is not directly related to the Red Hat exams. There is no evidence from the Red Hat Exam Prep guide or course outlines that you need to know how to download and write the RHEL 5 CDs during either exam.

Nevertheless, you need some Red Hat-style Linux distribution to prepare for the Red Hat exams, and this section focuses on whatever steps you need to take to download the relevant media. Obviously, it's best if you have an official subscription to the Red Hat Network, as you may get support for your installation and configuration. However, subscriptions are expensive, and courtesy of the GPL, alternatives are available.

The main alternatives are Fedora Core (versions 5 and 6), as well as the third-party "rebuilds" described earlier. Fedora Core 5/6 is the "testbed" for RHEL 5. While Fedora Core 5 is the "pre-alpha," Fedora Core 6 was used to develop the beta for RHEL 5. In other words, it's what Red Hat used to test many of the features now seen in RHEL 5. In my opinion, the better alternative is one of the rebuild distributions. As described earlier, these distributions are based on the source code for RHEL 5, which Red Hat has released under the GPL.

Several of these groups provide regular updates. As Red Hat releases updates to RHEL 5, these groups rebuild the updated source code into repositories. Most are compatible with the yum update tool. (Starting with release 7, Red Hat has dropped the "Core" from the name of the Fedora distribution.) You can keep these rebuild distributions up to date using yum or another update tool such as **smart** or **apt**.

Downloading Red Hat Enterprise Linux

Naturally, if you have a subscription to the Red Hat Network, it's best if you download RHEL 5 CDs directly from the network. Of course, you can purchase a subscription from Red Hat. But if you do not want to spend the money, Red Hat offers a trial subscription to RHEL 5. For more information, see www.redhat.com/rhel/details/eval/. As of this writing, it requires a quick telephone call from a Red Hat sales representative. If approved, you should be able to create a trial account on the Red Hat Network, from where you can download the installation CDs for RHEL 5 and receive updates for the 30-day evaluation period (though Red Hat may revoke this offer at any time). Once you have a Red Hat Network account, navigate to https://rhn.redhat.com, log into your account, click the Channels link on the top bar, and click the link associated with the following statement: *download ISO images of channel content*. (As Red Hat updates its systems, these URLs and links are subject to change.)

With an authorized Red Hat Network account, you should then be able to download the installation CDs for RHEL 5 in ISO format, using the instructions therein.

As of this writing, Red Hat Enterprise Linux is not available on DVDs, but I believe will be sometime in the near future.

Red Hat Enterprise Linux Source RPMs

As of this writing, you can navigate to the Red Hat FTP server, log in anonymously, and download the source code associated with RHEL 5. You can use commands such as rpmbuild to build the source code packages into binary RPMs that can then be installed.

Unfortunately, it is rather difficult to take these steps with all RHEL 5 source RPMs. Fortunately, a number of third parties have "rebuilt" these source RPMs into working distributions.

Third-Party Rebuilds

When third parties rebuild the RHEL 5 source RPMs, they do so under the GPL. However, they still have to respect various trademark laws that prohibit copying without permission. So when third parties rebuild RHEL 5 packages, they create their own icons, logos, and backgrounds.

However, several rebuild distributions have done more. For example, CentOS (www.centos.org) include installation DVDs and even live CDs similar to those associated with the Ubuntu and Knoppix distributions. They may even be available at the same servers from which you can download other Linux distributions. For example, the CentOS list of mirrors for their (and I say their, because CentOS is a community of developers) rebuild is available from www.centos.org/modules/tinycontent/index.php?id=13.

The Fedora Core 5/6 Prep Option

It's possible to prepare for the Red Hat exams using Fedora Core 5 or 6. Red Hat developed RHEL 5 from this distribution. Fedora Core 5/6 is freely available online and is easily downloadable. It's also available from many other sources, including my book *Mastering Fedora Core 5*, published by Sybex.

However, Fedora Core 5 is essentially a pre-alpha version of RHEL 5, and as such it may not reflect the RHEL 5 you work with during the Red Hat exams or on the job. While Fedora Core 6 is a bit closer, the software packages will vary. For that reason, I believe if you can't afford a subscription to the Red Hat Network, the best option is one of the third-party rebuild distributions.

An Overview of the Download Process

Whether you download RHEL 5, a rebuild, or Fedora Core 5/6, the basic download process is the same and follows these basic steps:

1. Select a distribution to download.
2. Find the download server with ISO files.
3. Proceed with the download, using a high-speed connection.

Other options are possible; for example, you can install Fedora Core 5/6 or some of the rebuild distributions directly from their Internet servers. As Red Hat supports downloads from HTTP and FTP servers, all you need is a boot disk and a sufficiently high-speed Internet connection. But when you download installation media, you can use that media to install RHEL 5 again and again on multiple computers.

You can then use a command such as **cdrecord**, or a GUI tool such as GnomeBaker, K3b, or even many Microsoft Windows-based tools to write the ISO file to appropriate blank CDs. The use of GUI tools to write ISO files to CDs (or even DVDs) is fairly trivial. Just look for the menu command that writes the ISO directly to the CD or DVD.

Downloads are not practical without a high-speed connection. (I once tried downloading a Red Hat CD over a telephone modem. After three days, the download file was corrupt and unusable.) If you don't have a high-speed connection, RHEL 5 CDs are available from Red Hat (though they're expensive), or CDs associated with some of the rebuild distributions are available from third parties. For example, the rebuilds created by CentOS and Scientific Linux (www.scientificlinux.org) are available for a modest fee from CheapBytes (www.cheapbytes.com).

Certification Summary

The RHCE and RHCT exams are not for beginners. This chapter covers the prerequisites for the RHCE exam and thus the elementary skills that you need for the remainder of this book. If the explanations in this chapter are too brief, you may need to refer to sources such as those I cite at the beginning of this chapter. While these exams are based on RHEL 5, you can use Fedora Core 5/6 or a third-party rebuild of RHEL 5 to study for these exams.

This chapter provides an overview of many Linux fundamentals. While the RHCE and RHCT hands-on exams may not explicitly test the skills you learn in this chapter, you need to know many of these fundamentals to solve the problems presented on those exams.

Before you start planning your Linux installation, you need a basic degree of knowledge of PC hardware, specifically the Intel-based architecture. A basic understanding of IRQ ports, I/O addresses, DMA channels, and hard drive systems can help you plan how Linux manages and connects every component in your PC.

But not all hardware is supported by Linux. You should now have enough information to find the hardware that fits your needs. Alternatively, you now know about the resources that help you determine what other hardware you need that also works with Linux. Planning your Linux installation makes it easier to handle a wide variety of hardware.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 1.

Basic Hardware Knowledge

? The Red Hat exams are given on computers built for an Intel-based 32-bit architecture.

? An Intel-architecture PC has three basic communications channels: IRQ ports, I/O addresses, and DMA channels.

? The latest version of Red Hat Enterprise Linux as certified should have at least 256MB of RAM.

? You can set up Linux on IDE, SCSI, USB, or IEEE 1394 hard drives. However, the BIOS of a PC can load Linux boot files only from the first two PATA, SATA, or SCSI drives.

Basic Linux Knowledge

? Linux is managed through a series of text configuration files.

? Understanding text editors is a critical skill. If you ever have to recover your system with a rescue CD, you may not have access to the GUI and will need to know how to use a console-based text editor such as vi.

Linux Filesystem Hierarchy and Structure

? Linux directories are organized to the Filesystem Hierarchy Standard (FHS).

? In the FHS, devices such as mice and hard drives are grouped in the /dev directory. Some /dev files have logical names such as dvdwriter and are linked to the actual device files.

? FHS partitions can be managed and formatted with the fdisk, fsck, and mkfs commands.

? The Logical Volume Manager allows you to consolidate multiple partitions in one filesystem, on one directory.

? Once configured, Linux directories can be mounted on a partition through /etc/fstab or directly with the mount command.

Basic File Operations and Manipulation

? Linux administrators need to know how to use the command line interface.

? Basic commands allow you to navigate, find the files that you need, read file contents, create new files, and more.

? File filters allow you to search through the files themselves for specific citations or other file characteristics.

? Administrative commands allow you to manage Linux in a number of ways, including running processes and managing logged-in users.

Printing

? The default Red Hat Enterprise Linux print system is CUPS.

? You can configure printers by directly editing the files in the /etc/cups directory or by opening the Red Hat Printer Configuration tool with the **system-config-printer** command.

Shells

? Command lines are based on a shell.

? With the right permissions, you can set up shell programs in executable scripts.

? The way a shell works depends on the settings in its variables and parameters. Some variables and parameters are grouped in the inherited environment, which maintains settings from shell to shell.

? With stdin, stdout, and stderr, you can manage different data streams.

Basic Security

? Basic security within Linux is based on file permissions, users, groups, and umask.

? The SUID and SGID bits allow you to share owner-level permissions with different users and groups.

? Shadow passwords hide user authentication data. The Shadow Password Suite protects user and group passwords in files that should be accessible only to the root user.

System Administration

- ? While it's normally best to log in as a regular user, it's faster to log in as the root user for the RHCE and RHCT exams.
- ? Standard files for new users are kept in /etc/skel.
- ? Daemons are processes that run in the background.
- ? Network service can be controlled through scripts in the /etc/init.d and /etc/ xinetd.d directories.
- ? The cron daemon helps you schedule different jobs, including backup and restore jobs, which should be done when network use is at a minimum.
- ? When you have problems, system log files, as organized by /etc/syslog.conf, provide important clues to the causes.

Basic TCP/IP Networking

- ? Most of the work in TCP/IP networking is in configuring IP addresses.
- ? There are three different sets of private IPv4 addresses suitable for setting up TCP/IP on a LAN.
- ? IPv6 addresses include all available IPv4 addresses. If the first three bits of an IPv6 address are 001, that is a unicast address-in other words, one that is associated with a specific computer or other device.
- ? The first 48 bits of an IPv6 address are typically associated with a specific network.
- ? Tools such as **ping**, **ping6**, ifconfig, and netstat can help you diagnose problems on that LAN.
- ? Name resolution configuration files determine how your computer finds the right IP address.

Familiarity with Standard Network Services

- ? There are a number of standard network services, including NFS, sendmail, POP, IMAP, FTP, DNS, DHCP, Samba, Apache, and NIS.
- ? Each of these services, when installed, can be configured to start and stop through the scripts located in the /etc/rc.d/init.d or /etc/xinetd.d directories.

Basic Network Security

? Basic network security settings can depend on allowing or denying access to different computers by their IP addresses or by the desired TCP/IP port.

? Computers behind a firewall can be protected through Network Address Translation or various iptables commands.

Other Basic Prerequisite Skills per the Red Hat Exam Prep Guide

? While GUI e-mail clients should be trivial, it's important to know how to configure a command line e-mail client.

? While GUI Web browsers should be trivial for serious Red Hat exam candidates, it can help to know a text-based browser such as **elinks**.

? While GUI FTP clients should be trivial for serious Red Hat exam candidates, it can help to understand a text-based FTP client such as lftp.

Downloading the Red Hat Enterprise Linux Installation CDs

? There is no evidence that you need to know how to download the Red Hat installation CDs for the Red Hat exams.

? While the best option is to download the RHEL 5 CDs from the Red Hat Network, excellent options are available.

? You can use the rebuild distributions to prepare for the Red Hat exams. Their distributions are built on the same source code used by Red Hat for RHEL 5.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As there are no multiple choice questions on the Red Hat exams, there are no multiple choice questions in this book. These questions exclusively test your understanding of the chapter. While the topics in this chapter are "prerequisites," it is okay if you know another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer for many of these questions.

Basic Hardware Knowledge

1\.

Once Linux is booted, what file can tell you all about the CPU(s) on your system?

____

Basic Linux Knowledge

2. If you're editing the /etc/inittab file in vi, what command would you use to copy the currently highlighted line? ____ ?

Answers

2\. The **yy** command in vi yanks a copy of the current line into the buffer.

Linux Filesystem Hierarchy and Structure

3. What command would you use to find currently mounted drives? ________________ ?

Answers

3\. To find currently mounted drives, the simplest method is to use the mount command; other commands that can tell you about mounted drives are **cat /etc/mtab** and **df**.

Basic File Operations and Manipulation

4. If you want to find the actual number of times user mj is logged into your Linux computer, what command would you use? __________________________________ ?

Answers

4\. The most elegant way to find the actual number of times user mj is logged in is to use the **who | grep mj | wc -l** command. However, for these purposes, it's usually good enough to count from the list associated with the **who** command. Remember that results are what matter on the Red Hat exams.

Printing

5. You're maintaining a large queue of print jobs on your network, and you need some job numbers to make sure the engineers get highest priority on the printer. What command would you use to list print job numbers? ___________________________________

Answers

5. The simplest method to check print queues is to use the lpq command. If you want more information, try **lpq -l** for a long listing format. If you have more than one printer, the **lpq -a** command checks all configured printers.

Shells

6. What command would you use to add the /usr/sbin directory to your PATH?

____

Answers

6. The simplest way to add /usr/sbin to your $PATH is to use the **PATH=$PATH:/usr/sbin** command. But to make sure this command takes effect the next time you log in, you should add this command to the hidden .bash_profile file in your home directory.

Basic Security

7. When you run the umask command, you see the following result: 0000. The next time you create a file, what will be the permissions? ___________________________

Answers

7. The answer is read and write permissions for all users, or _rw_rw_rw. Even if you try to set it to allow execute permissions, Red Hat won't let you do this anymore. You'll need to set execute permissions on each file after creation.

System Administration

8. Based on the following line from a user's crontab file, when will the Berkeleylives program be run?

```
0
01  23  * Berkeleylives
```

Answers

8. The convention for the first five entries in a crontab line is minute, hour, day of month, month, and day of week, so this particular job will run at 1 A.M. on March 2.

Basic TCP/IP Networking

9. Provide an example of an appropriate IPv4 network address, subnet mask, and broadcast address for a network of less than 300 computers on the 10.0.0.0 private network. ?

Network address: ________________________

Subnet mask: ________________________

Broadcast address: ________________________

Answers

9. The answer to this question will vary widely. If you can't provide many answers to this question, you may need to learn more about basic IPv4 addressing. One example is a network address of 10.11.12.0, a subnet mask of 255.255.255.0, and a broadcast address of 10.11.12.255.

Familiarity with Standard Network Services

10. What is the protocol associated with the service used to connect Linux to a Microsoft Windows-based network? ____ ?

Answers

10. The basic protocol is Samba (SMB), associated with the Common Internet File System (CIFS). Any of these answers are correct; however, as described in Chapter 10, it's important to know that the CIFS module has superseded the SMBFS module for RHEL 5.

Basic Network Security

11. What command would you use to start the basic Red Hat Enterprise Linux firewall configuration utility? ?

____

Answers

11. The basic Red Hat firewall configuration utility can be started with the **system-config-securitylevel** command.

Other Basic Prerequisite Skills per the Red Hat Exam Prep Guide

12. Name one of the major e-mail clients available on Red Hat Enterprise Linux 5. ?

____

Answers

12. Several different e-mail clients are available in RHEL 5, including mutt, mail, Kmail, and Evolution.

Downloading the Red Hat Enterprise Linux Installation CDs

13. Name one alternative distribution that uses the same source code as RHEL 5. ?

____

13. There are several different rebuild distributions available, including CentOS, Lineox, Scientific Linux, and more. One list of available rebuild distributions can be found at http://linuxmafia.com/faq/RedHat/rhel-forks.html. The answer is correct only if the group you've cited has rebuilt the source code for RHEL 5.

PREV NEXT

Lab Questions

The first lab is fairly elementary, designed to get you thinking in terms of networks and networking. The last two labs both work with the /etc/inittab configuration file. Before working with that file, make sure to back it up first.

Lab 1

1. You have 18 computers on a LAN behind a firewall. Diagram your computers on a sheet of paper. Connect them together in a "star" configuration. Assign a private IP address to each computer. Take one computer and draw a second connection to the Internet. ?

While this is a fairly simple exercise, Linux is built for networking. To understand what you can do with Red Hat Enterprise Linux, you need to think in terms of the role of your computer on a network.

Answers

1. There are many ways to configure the IP addresses on a LAN. But it is generally best to do it by setting up a network from one of the private IP address ranges. When you configure networking on your LAN, pay particular attention to the computer that also has a connection to the Internet. The IP address of its connection to your network will be the gateway address for every other computer on your LAN. It's also the logical location for any firewall that you may wish to configure.

Lab 2

2. In this lab, you'll start your experiments with the /etc/inittab file. So before you begin, back it up to a file such as /etc/inittab.bak or back up a copy to your home directory. ?

1. Use the vi editor to open the /etc/inittab file in your computer.
2. Take a look at your id variable. If it's set to 3, change it to 5; if it's 5 set it to 3.
3. Reboot your computer and see what happens.
4. Restore your original /etc/inittab file.

Answers

2. When you troubleshoot a Red Hat Enterprise Linux computer, one of the things you'll be checking are critical configuration files. One key file in the boot process is /etc/inittab. One thing that I can do in this book is to illustrate the behavior of potential problems. The more problems you're familiar with, the easier it will be to troubleshoot or debug a problem during the RHCT and RHCE exams. However, there is often more than one way to solve a problem. I present one method, but you may be able to find others.

   To go through this lab, I'd take the following steps:

   1. Log in as the root user. You can do this from either the GUI or the text login interface.
   2. Run the **cp /etc/inittab /root/inittab** command. This backs up the subject configuration file in the root user's home directory.
   3. Open the subject file with the **vi /etc/inittab** command.
   4. Scroll down to until you see the following line:

Lab 3

3. In this lab, you'll experiment a bit more with the /etc/inittab configuration file.

   1. If you haven't already done so, create a backup for /etc/inittab.
   2. Press CTRL-ALT-F2. You should see a virtual console text login screen.
   3. Return to the original text console by pressing CTRL-ALT-F1 or the GUI console by pressing CTRL-ALT-F7.
   4. In the /etc/inittab file, identify the lines related to the virtual login consoles.
   5. Try experimenting with these lines with the **mingetty** commands. Add a comment character (#) in front of the second line with the **mingetty** command.
   6. Run the **init q** command to make Linux reread this file.
   7. Try pressing CTRL-ALT-F2 again. What happens?
   8. Restore your original /etc/inittab configuration file.

Answers

3. In this lab, experiment with deactivating a specific virtual console. By default, six virtual text login consoles are configured in the /etc/inittab configuration file. You'll deactivate the second of the six consoles.

   1. Log in as the root user. You can do this from either the GUI or the text login interface. If you're in the GUI, open a text console. Right-click the desktop and choose New Terminal in the pop-up menu.
   2. Run the **cp /etc/inittab /root/inittab** command. This backs up the subject configuration file in the root user's home directory.
   3. Open the subject file with the **vi /etc/inittab** command.
   4. Scroll down until you see the following line:

1. The key Linux file associated with the CPU is /proc/cpuinfo. However, other files can tell you about how Linux detected the CPU, including /var/log/dmesg. The more you know about such files, the more problems you can diagnose with your hardware.

Chapter 2: Hardware and Installation

Overview

Installation is one of the two parts of both the RHCE and the RHCT exams. To pass this part of each exam, you'll need to know a lot more than just the basic GUI installation process for a single computer! Once you've studied the installation chapters (Chapters 2 and 5), you'll be able to install Red Hat Enterprise Linux (RHEL) in a number of ways: over a network, directly from the CD, using boot disks, and with automated kickstart-based tools.

While this chapter covers the "basics," they are important. Naturally, you need compatible hardware, as well as sufficient processing power and RAM. The latest Linux tools help detect portable and network devices.

Generally, unless you've copied the installation DVD or CDs to the local system, network installations are fastest. To help you learn all installation methods, we'll show you how to configure a network installation server.

On the Job

While Red Hat does not provide an installation DVD, it is available for some of the "rebuild" distributions such as CentOS and Scientific Linux.

Then you'll learn how to install RHEL on your system, step by step. The installation program is known as Anaconda. You'll see how you can create regular (software) RAID and Logical Volume Management (LVM)-based partitions. Then you'll be guided through the First Boot process to see how to continue the RHEL configuration process after installation is complete.

You may be asked to install and configure some or all of the services described on the Red Hat Exam Prep guide (www.redhat.com/training/rhce/examprep.html). This chapter will help you understand the services that you can install with RHEL. The fastest way to install RHEL is in text mode.

However, you may wish to install RHEL via the regular graphical screen as you can't customize LVM partitions in text mode. If the test system and network is up to date (for example, Fast Ethernet is probably sufficient), the time penalty associated with GUI installations may be trivial.

Certification Objective 2.01-Hardware Compatibility

Now it's time to explore the hardware that RHEL can handle. While some manufacturers now include their own Linux hardware drivers, most Linux hardware support comes from third parties. Fortunately, a vast community of Linux users are hard at work, producing Linux drivers and more, even distributing them freely on the Internet. If a certain piece of hardware is popular, you can be certain that Linux support for that hardware will pop up somewhere on the Internet and will be incorporated into various Linux distributions, including RHEL.

Inside the Exam

Focus During Installation

Both the RHCE and RHCT exams include an Installation and Configuration section. You'll have to do more than just install Linux. You'll follow a series of instructions, configure custom partitions, and configure certain services.

Time limits are severe on these exams. Install and configure as much as you can when you install RHEL on your computer. Although you can configure and install almost anything after Linux is installed, that can take more time than you have.

On the other hand, don't install everything. It takes time to install gigabytes of software over a network. If you're spending time installing software that you don't need, that's time you can't get back during the exam.

As you read this chapter, learn every part of the installation process. Know what you need to install. For example, if you see a requirement to set up Apache and Samba servers, you'll want to install the Web Server and Windows File Server package groups when you install RHEL.

Studying for the Installation and Configuration Section

You can use one of the "rebuilds" of RHEL 5 or even Fedora Core 6 to study for the Installation portion of the RHCE and RHCT exams. The steps required are essentially identical to those for RHEL 5. To assure you that the steps are the same, consult "Installing Red Hat Enterprise Linux 5" in the Online Learning Center (http://highered.mhhe.com/sites/0072264543), which provides a pictorial-only guide to the RHEL 5 installation process.

Be careful when purchasing a new computer to use with Linux. Though Linux has come a long way the last few years, and you should have little problem installing it on most modern PCs, you shouldn't assume Linux will install or run flawlessly on *any* PC, especially if the PC in question is a state-of-the-art laptop computer (though several major laptop manufacturers seem determined to maintain good relationships with the Linux community). Laptops are often designed with proprietary configurations that work with Linux only after some reverse-engineering. For example, when I installed Red Hat Enterprise Linux on a new widescreen laptop, I could install only in text mode.

Other types of hardware, such as "winmodems" and "winprinters," are designed to use Microsoft Windows driver libraries. Integrated hardware (such as video chips that share system RAM) and parallel port devices can also be problematic. While Linux drivers exist for many of these devices, do your research.

Linux runs very well on lower-end computers. This is one of Linux's strong points over other operating systems, especially Microsoft Windows Vista. Linux runs fine on 64MB of RAM, although more is always better, especially if you want to run any graphical applications. RHEL 5 does require a minimum of 192MB of memory to start the graphical installer. However, the latest versions of Linux do have limits; modern distributions don't run on anything less

than a Pentium-class system.

Exam Watch

While it is important that you know how to select and configure hardware components to get to a smoothly running Linux computer, the RHCE and RHCT exams are not hardware exams.

Linux Hardware Documentation

Many resources are available to help you select the best hardware for Linux. Thousands of Linux gurus are available online via mailing lists, IRC rooms, and newsgroups. Perhaps the best places to look are the Linux Documentation Project (LDP) or the Red Hat Hardware Compatibility List (HCL). The LDP is a global effort to produce reliable documentation for all aspects of the Linux operating system, including hardware compatibility.

Linux Hardware HOWTO

The Linux Hardware HOWTO is a document listing most of the hardware components supported by Linux. It's updated irregularly with added hardware support, so it is a relatively up-to-date source of information, available at www.tldp.org.

The Red Hat Hardware Compatibility List

The Red Hat HCL is different from the one you'll find in the Linux Hardware HOWTO. It specifies name brand hardware that has been tested with various versions of RHEL. If you've purchased RHEL, Red Hat will provide some level of installation support for any certified or compatible hardware. Some hardware that has been tested by Red Hat has specifically been found not to work with Red Hat Linux or RHEL and is therefore not supported. Red Hat doesn't have the resources to test more than a limited range of hardware; most PCs and servers built today work well with RHEL. For that information, refer to the aforementioned Linux Hardware HOWTO.

Plug and Play and the Hardware Abstraction Layer

Plug and play (PnP) refers to the ability of an operating system to allocate hardware ports or addresses automatically to specific devices such as hard drives, sound cards, or modems. Linux's ability to work with PnP devices is finally up to speed, courtesy of the Linux implementation of the Hardware Abstraction Layer (HAL). Conceptually different from the Microsoft version, HAL provides a constant list of detected components. Some distributions can now automatically detect and mount the smart cards associated with digital cameras and fingerprint readers.

If you want to see the full list of detected hardware, run the **lshal** command. It's a long list; you may need to pipe the output to a pager with a command like:

```
#
# lshal | less
```

Certification Objective 2.02-CPU and RAM

Red Hat Enterprise Linux supports computers with Intel and compatible 32-bit and 64-bit processors.

Linux is commonly used as a server operating system. Many server applications can take advantage of the flexibility provided by multiple CPUs. This is known as *symmetric multiprocessing (SMP)* support. Linux began supporting multiple CPUs with the release of the 2.4 kernel back in 2001. RHEL 5 also supports virtual machines with a customized Xen-based kernel. If you have a newer "dual-core" type CPU, RHEL 5 can even support hardware virtualization, which allows dedicated installations of Microsoft Windows (and other Intel-compatible operating systems) within Linux.

Compatible CPUs

You can install RHEL 5 on systems with a wide variety of CPUs. Red Hat supports six different CPU architectures:

- x86
- Athlon/AMD64 (x86_64)
- Itanium (ia64)
- IBM zSeries
- IBM iSeries
- IBM pSeries

Exam Watch

As of this writing, we assume that Red Hat tests and will continue to test based on the most popular architecture, x86. Other architectures such as Itanium use a different boot loader (ELILO), which is not covered in the Red Hat Exam Prep guide or associated course outlines.

On the Job

Some developers hope to increase the SMP limit to 128 CPUs. If you're running Linux on an SMP computer, keep up to date with the latest kernel developments at www.kernel.org.

CPUs and Virtualization

Red Hat is in the process of incorporating virtualization in its operating systems. Both Fedora Core 6 and RHEL 5 include Xen, which is a "free virtual machine monitor," which includes QEMU-based emulation to support virtualization similar to VMware.

On the Job

QEMU is one more alternative for virtualization, licensed under the GPL and the closely related Lesser GPL. Kernel-based Virtual Machine (KVM) technologies require Linux kernel version 2.6.19; RHEL 5 uses 2.6.18. Paravirtualization for VMWare isn't expected until 2.6.21.

There are two kinds of virtualization associated with Xen. Paravirtualization provides a software interface that allows you to install specially ported operating systems (with a Xen-enabled kernel) within software-based virtual machines. Full, or hardware-assisted, virtualization supports direct hardware access; it is limited to certain Intel Dual Core and AMD X2 CPUs.

On the Job

Not all Intel Dual Core or AMD X2 CPUs support hardware-assisted virtualization. AMD X2 CPUs need to be TL-50 and above; Intel Dual Core CPUs need to be T2300 (or T5600) and above. Furthermore, Intel-based systems are often disabled in the BIOS and may not support hardware-assisted virtualization unless specifically activated through the BIOS menu. If you see the vmx (Intel) or svm (AMD) flags in /proc/cpuinfo, your system supports hardware virtualization.

RAM Requirements

The minimum RAM requirements for RHEL 5 are trivial for today's computers. While you need at least 192MB to install in graphical mode, 64MB is sufficient to install in text mode-and to run this distribution with a text-based login. One advantage of this small footprint is that it allows you to configure more virtual machines using Xen.

On the Job

In reality, the minimum amount of RAM depends on the amount of shared video RAM. For example, if 64MB of RAM is used for your video system, you need at least 256MB of RAM to install RHEL 5 in graphical mode.

Certification Objective 2.03-Hotswap Buses

After a lot of work over the years, Linux handles hotswappable devices well. If everything works as it should, you can plug a device into a hotswap system, and Linux automatically detects the device, loads drivers, and, if appropriate, mounts the data from that device on an appropriate filesystem. There are several commands available to help manage these devices.

You can install many devices externally to your computer. These devices are sometimes known as peripherals, which fall into six categories: serial, parallel, USB, IEEE 1394, smart cards, and PC Cards. A device attached to a serial port, such as a mouse or a modem, uses the device associated with that port. Devices attached to parallel, smart cards, USB, or IEEE 1394 ports normally use their own device files. PC Cards are a special case normally associated with laptop computers.

While Linux normally recognizes basic devices attached to these ports, configuring a few devices may take additional work.

Serial Ports

In many cases, configuring a device for a serial port is as simple as linking to the driver of the associated port. For example, if you have an external modem connected to the only serial port on your computer, the Linux HAL subsystem may have already linked the device for that port with the device for your modem. Run the **ls -l /dev/modem** command. If it shows something like the following output, you know that Linux has already linked your modem driver with the second serial port:

```
l
lrwxrwxrwx  1  root  root  10 Apr 1 13:17 /dev/modem -> ttyS1
```

Certification Objective 2.04-Configuring a Network Installation

Most Linux users can install RHEL from a CD/DVD. During the installation portion of the Red Hat exam, you'll save time by installing RHEL over a network from an NFS, HTTP, or FTP server. As you'll want to practice with network installations, you should set up a network server. For completeness, while it's not a "network installation" per se, you can also install RHEL from ISO images on a local hard disk.

On the Job

SELinux is an excellent system that provides a different level of security for Linux. In this book, I describe many ways that you can work with SELinux enabled. However, users do report trouble making SELinux work with Linux and services such as NFS, HTTP, and FTP. While SELinux has just been included in the Red Hat Exam Prep guide, the exams don't test how you've configured a network installation server. So, if necessary, disable SELinux for a specific service, or disable it completely with the Security Level Configuration tool described in Chapter 15.

Configuring a Network Installation Server

Once you have the Red Hat installation media, configuring a network installation server is a fairly easy process. All you need to do is copy the files from each CD to a common directory, configure sharing on the directory, and then activate the NFS, FTP, or HTTP network. Naturally, if you've downloaded a rebuild installation DVD, the process is simpler.

Before you set up a network installation source, you'll need a partition with at least 2.7GB of free space (3.5 if you're installing the RHEL 5 Client, appropriate for the RHCT exam). I'll illustrate the process for an NFS server and explain the variations for FTP and HTTP servers.

Creating an NFS Installation Server

In the following steps, you'll learn how to create a shared directory, copy the Red Hat installation files, and then set up the share through NFS. As NFS is the most efficient way to share files between Linux and Unix computers, I suspect it's the most likely option for network installations during the exam. You'll need the Red Hat Enterprise Linux installation CDs, or at least the ISO files associated with those CDs.

1. Create a directory for your installation files. With the following command, create the /inst directory:

Certification Objective 2.05-The First Installation Steps

Before installing Red Hat Enterprise Linux, let's examine some critical decisions that you'll make during the Red Hat exams. Time is of the essence on the exams. While you could just install everything, that could easily cost you 15 minutes or more. On the other hand, while you can customize individual packages to be installed, that can also steal the time that you need to configure the critical services required to pass the exam. By the time you're done reading this chapter, you'll know how to select just what you need.

You have to answer many interrelated questions during installation, just as you have many ways to access installation files, and many options on how to install the operating system. The following installation outline is designed to get you through the process as simply as possible. While other sections and chapters address the special situations that you're more likely to encounter on the RHCE and RHCT exams, you need to know how to install Red Hat Enterprise Linux before you can work though the other installation scenarios.

Boot Options

During the RHCE or RHCT exams, you'll have access to the installation files. There are four methods available to start the RHEL Desktop or Server installation process:

- Boot from a copy of the Red Hat Enterprise Linux installation CDs or DVD
- Boot from the first RHEL installation CD or DVD
- Boot from a special RHEL boot CD or USB key
- Boot from a kickstart server using a PXE network boot card

The last three options generally assume that you're going to install RHEL over a network. It's possible that you'll see one of these options during the exam, which is why I described how to create a network installation server earlier in this chapter.

Exam Watch

During the exam, avoid installing RHEL from the CDs if at all possible. The basic installation takes longer, and you don't want to waste time removing and reinserting CDs.

Booting from the First CD/DVD

Most current Intel-based PC hardware systems allow you to boot directly from the CD drive. You can start the installation process by booting from the first Red Hat Enterprise Linux CD or DVD.

If You Need an Installation USB or CD/DVD

If you don't have the first installation CD/DVD, you can start a network installation from a specialized boot USB key

or CD. While one may be provided for you for the RHCE or RHCT exams, you'll need to know how to create one so you can practice for the exam.

It's easy to create an installation USB key or CD from one of the files on the /images directory on the first installation CD:

- **diskboot.img** For a boot USB key
- **boot.iso** For a boot CD

The boot.iso file is small enough to fit on a credit card-sized CD. It contains all the information in the diskboot.img file. Alternatively, if your systems have USB ports and can boot from that media, you can create the installation media from the diskboot.img file.

For the purpose of this section, assume you've inserted the first installation CD/DVD into its drive, and it's automatically mounted in the /media/disk directory. (In practice, the /media subdirectory is named after the label on the CD/DVD; you may end up with a directory such as /media/RHEL-5 i386 Disc 1.) If automounting does not work, you may need to mount the CD/DVD yourself.

Creating a Boot USB Key

You can also create images on a USB key with the **dd** command from any running Unix or Linux computer. Find a USB key, save anything important that you've stored onto it, and insert it into your system. Run the **fdisk -l** command to find the device associated with the USB key. Assuming it's /dev/sdc, run the following commands:

```
#
# dd if=/media/disk/images/diskboot.img of=/dev/sdc
```

Certification Objective 2.06-Configuring Partitions, RAID, and LVM

A disk drive requires a partition table. The *partition* is a logical sequence of cylinders on the disk, while a *cylinder* represents all the sectors that can be read by all heads with one movement of the arm that contains all these heads. Although it's possible to create more, RHEL will recognize only up to 16 partitions on any individual SATA/SCSI or an ATA/IDE hard drive. But don't be too concerned about these details; you probably won't have to create so many partitions on either exam.

Once you create a partition, you can mount a directory directly on that partition. Alternatively, you can designate that partition as a RAID device or as part of a logical volume. Both systems are described briefly in this chapter; for more information, see Chapter 8.

On the Job

The main Linux partition utilities are fdisk and *parted*. As you'll see in Chapter 4, although you can create more than 16 partitions on each physical hard disk, you'll find that you can't write more than 16 partitions on each hard disk.

RAID, Briefly

RAID is short for a Redundant Array of Independent (or Inexpensive) Devices, and it can help you create a filesystem on more than one partition, with or without redundancy. Red Hat supports the software version of this scheme at four different levels.

- RAID 0 stripes a filesystem across multiple partitions for faster reads and writes.
- RAID 1 mirrors all data in a filesystem between two partitions.
- RAID 5 uses parity bits (striping) for data redundancy over three or more partitions.
- RAID 6 uses double parity bits (striping) for data redundancy over four or more partitions.

To take full advantage of RAID, you'll want to configure partitions on different physical hard drives, preferably connected to different hard drive controllers. That maximizes the ability to read and write to a RAID filesystem, and it means (except for RAID 0) that your data will not be lost if one hard drive fails. But multiple hard drives may not be available on the computer you use to take an exam. (For more information on RAID and how you can manage it after installation, see Chapter 8.)

Logical Volumes, Briefly

Logical volumes give you flexibility with your filesystems. You can expand and contract the space associated with different directories. If you need more space, say for the /home directory, you can add a new hard drive, partition it as a logical volume, and use it to add more space to the /home directory.

The Red Hat installation process, by default, configures all but the /boot directory as part of a logical volume, with the configuration described in Table 2-1. Chances are good that won't match the requirements on your RHCE or RHCT installation exam. We'll explore how you can customize it shortly.

Table 2-1: Result When You Partition Automatically

Location	**Size**
/boot	100MB
Swap	Twice available RAM (assuming sufficient hard drive space)
/	Remaining space on the drive, mounted on a Logical Volume Group

It's faster and easier to configure LVM during the installation process. While Red Hat provides the Logical Volume Management GUI tool, it takes time to use it. It may be faster to use regular command line-based commands to reconfigure LVM. But requirements change and mistakes are made, during the installation process and in real life. So for more information on logical volumes and how you can manage them after installation, see Chapter 8.

Naming Conventions

Linux has a simple naming standard for disk partitions: three letters followed by a number. The first letter identifies the type of drive (*h* is for IDE/EIDE, *s* is for SATA or SCSI). The second letter is *d* for disk, and the third letter represents the relative position of that disk, starting with *a*. In other words, the first ATA/IDE drive is *hda*, followed by *hdb*, *hdc*, and *hdd*.

The number that follows is based on the relative position of the primary, extended, or logical partition. Primary partitions can contain the boot files for an operating system. Hard drives can also be configured with one extended partition, which can then contain up to 12 logical partitions.

You are limited to four primary partitions on each hard disk. But four partitions are often not enough. If you need more partitions on a hard drive, substitute an extended partition for one primary partition. You can then configure the logical partitions that you need within the extended partition.

You can't install files directly in an extended partition. You must first allocate some extended partition space to at least one logical partition. You can then configure logical partitions within that extended partition. In all cases, the first logical partition on the first PATA/IDE drive is hda5.

Each partition is associated with a Linux device file. At least this is straightforward; for example, the device filename associated with the first logical partition on the first PATA/IDE drive is /dev/hda5.

Exam Watch

You should know the device name associated with each partition, as well as the starting names and numbers of any logical partitions created on any basic disk drive. Also remember that logical partitions on a hard drive always start with number 5; on the first PATA/IDE hard drive on a PC, that is hda5.

Exercise 2-1: Partitioning

You may never have had to plan partitions on a basic Microsoft Windows desktop computer. On a real server, whether you're using Windows or Linux, you should preplan your disk usage and partitions very carefully. This is a preliminary exercise; be prepared to think more deeply about partitions later in this chapter and in Chapter 4.

1.

On a piece of paper, draw a rectangle to represent each hard drive on your computer.

2.

Label them in order just as Linux would (Hard Drive 1: /dev/hda, Hard Drive 2: /dev/sda, Hard Drive 3: /dev/sdb).

3.

Use this diagram to plan how you are going to partition each drive. While this is a preview of future chapters, you should already know that Linux is set up in multiple directories. Each of these directories can be set up in its own partition. Think about how much space you want to allocate to several major directories, such as /home, /var, /usr, /boot. Don't forget to allocate some area for a swap partition.

Using this method, you can organize your data, keeping system or users' files together, as well as strategically plan where to place your swap partition(s). Now in the second exercise, let's examine how you can create and configure partitions during the installation process. We'll also examine how you can allocate a filesystem to a partition, a logical volume, or a RAID array.

Exercise 2-2: Partitioning During Installation

To follow along in this chapter, you presumably are installing RHEL 5 on some system. It's easiest if you practice using a virtual machine based on VMware or Xen. If you make a mistake, you can restart the installation process and return to these partitioning steps. This exercise starts with Figure 2-8, which is the first partitioning step, and assumes you're testing with the graphical installation.

Figure 2-8: Basic partitioning

The Advanced Storage Configuration option is associated with network connections using the TCP/IP iSCSI network protocol. As this option is not part of the Red Hat Exam Prep guide, we won't cover it here. For more information, start with the Linux-iSCSI project, described at http://linux-iscsi.sourceforge.net.

1.

Click the top drop-down text box, and you'll see four options. In the following steps, you'll examine each option in turn with the associated default partitioning layout.

2.

Check the Review And Modify Partitioning Layout checkbox near the bottom of the window.

3.

Select Remove All Partitions On Selected Drives And Create Default Layout, and click Next.

4.

If more than one drive is available, deselect drives where you want to save the data.

5.

If you see the warning about removing partitions from specified drives, click Yes to continue.

6.

Review the default partitioning layout, and click Back to return.

7.

Select Remove Linux Partitions On Selected Drives And Create Default Layout, and click Next to continue. Repeat steps 4 to 6.

8.

Select Use Free Space On Selected Drives And Create Default Layout, and click Next to continue. Repeat steps 4 to 6.

9.

Select Create Custom Layout, and click Next to continue.

10.

Start creating your own custom layout. If you're starting with blank hard disks, no partitions will be configured. Delete configured partitions if no space is available.

11.

Try creating a regular partition. Click New. Create an appropriate mount point, such as /home/user. Click the File System Type drop-down text box and review the available formats.

12.

Create one partition with an unused filesystem type, and click OK to continue.

13.

Repeat steps 11 and 12. For the purpose of this exercise, the default 100MB is sufficient.

14.

Create five RAID partitions. Click RAID; this opens the RAID Options window. If you don't have any other RAID partitions, only one of the three options can be selected: Create A Software RAID Partition. Click OK; this opens the Add Partition window, with a Software RAID File System Type. Create a partition of the desired size (100MB is OK for this exercise) and click OK. Repeat this process three times, making sure to select Create A Software RAID Partition each time.

15.

You should now have five RAID partitions. Click RAID, and in the RAID Options window, select Create A RAID Device. Click OK to open the Make RAID Device window.

16.

Create an appropriate mount point such as /home/raidtest. Check available filesystem types; you'll see the same format options you saw in step 11.

17.

The RAID Device drop-down text box is trivial; it just specifies the device file associated with the RAID array you're about to create.

18.

Click the RAID Level drop-down text box. You'll see that you can configure four different types of RAID

arrays: RAID0, RAID1, RAID5, and RAID6. (For more information on each level, see Chapter 8.) Select RAID6.

19. Enter **1** in Number of Spares, and click OK. This works because RAID 6 requires a minimum of four RAID partitions and as many spares as available. Review the resulting RAID device. How much space is available in the device? How does that compare to the space used by the five different RAID partitions?

20. Now click New again, and create a partition of the Physical Volume (LVM) filesystem type. Repeat the process to create a second LVM partition.

21. Click LVM. This opens the Make LVM Volume Group window. Review the available options.

22. In the Make LVM Volume Group window, click the Physical Extent drop-down text box. Review the available Physical Extents, which are units associated with volume groups.

23. Make sure all available Physical Volumes To Choose are active.

24. Click Add; this opens the Make Logical Volume window.

25. Create an appropriate mount point such as /home/volume. Note that the available filesystem types are more limited. The Logical Volume Name shown is just the default; you can use any legal filename for your logical volume. Set a size that does *not* use all available space. Click OK.

26. Review the result in the Make LVM Volume Group window. Click OK and review the result in the original partition window.

27. Changes are not permanent; you should be able to click Back a couple of times and return to restore the original partition configuration.

Separate Filesystems

Normally, you should create several partitions when preparing your hard drive to install Linux. This is a good idea for various reasons. First, RHEL is normally configured with at least two filesystems: a Linux native filesystem and a Linux swap filesystem. Second, if you want to install RHEL and another operating system on the same computer, you should configure separate partitions for each operating system. You can and should configure software RAID partitions on different physical hard drives (if available) during the RHEL installation process. However, if you have a hardware RAID system, you'll need to configure it after RHEL is installed.

Exam Watch

During the Installation and Configuration portion of either exam, pay careful attention to the instructions. Make sure that the partitions you create while installing RHEL match any instructions you see. It's much more difficult and much more time-consuming to revise partitions with utilities such as fdisk or *parted* after RHEL is installed.

Stability and Security

Linux is organized in a Filesystem Hierarchy Standard (FHS) that includes a number of directories described in Chapter 1. You can organize these directories into a few or many hard drive partitions. During the installation process, RHEL is by default organized into three partitions: the root directory /, the /boot directory, and a swap partition. One recommended configuration for a Linux server includes separate partitions for each of the following directories: /, /boot, /usr, /tmp, /var, and /home. Other partitions may be appropriate for corporate data, database services, and even directories associated with Web (/var/www/html) and FTP (/var/ftp/pub) sites if you need them to hold a lot of data.

Partitioning the hard drive in this manner keeps system, application, and user files isolated from each other. This helps protect the disk space used by the Linux kernel and various applications. Files cannot grow across partitions. For example, an application such as a Web server that uses huge amounts of disk space can't crowd out space needed by the Linux kernel. Another advantage is that if a bad spot develops on the hard drive, the risk to your data is reduced, as is recovery time. Stability is improved.

Security is also improved. Multiple partitions give you the ability to set up certain directories as read-only filesystems. For example, if there is no reason for any user (including root) to write to the /usr directory, mounting that partition as read-only will help protect those files from tampering.

While there are many advantages to creating more disk partitions, it isn't always the best solution. When hard drive space is limited, the number of partitions should be kept to a minimum. For example, if you have a 4GB hard drive and want to install 3000MB of packages during RHEL installation, you may not want to dedicate extra space to the /var directory. You need room for swap space, additional programs, and your own personal files on other directories.

Exam Watch

It can take considerable time to set up LVM partitions. Unless you know the process very well, the fastest way is through the RHEL installation program in graphical mode, which is available when installing from CD/DVD or over a network from an NFS server. (LVM configuration is not available via text mode RHEL installation.) Learn the process well, just in case you need to set up LVM during the Installation and Configuration part of your exam.

Basic Storage Space Requirements

Linux is a flexible operating system. While a full installation of RHEL requires several gigabytes of space, slightly older versions of Linux fit even on a 1.44MB floppy disk. Depending on your needs, you can install RHEL, with a couple of services, *without the GUI*, on any hard drive larger than 2GB.

On the Job

There are also complete Linux distributions that you can boot and load directly from a CD or DVD, which can be used to diagnose hard disk failures on Microsoft Windows PCs. For more information, see www.knoppix.net. Even Fedora Core now has a so-called "live" DVD. But you won't be able to use live CDs or DVDs on the Red Hat exams.

You should size your Linux partitions according to your needs and the function of the computer. For example, a mail server will require more space in /var, because mail files are stored in /var/spool/mail. You could create a separate partition for /var or even /var/spool/mail. In almost every case, it's a good idea to configure at least the /boot directory on a separate partition.

On the other hand, if you install everything, including support for various languages, you could require 10GB or more.

Example: File Server

If the Linux system you are installing is to be a file server, then you could configure your partitions as shown in Table 2-2.

Table 2-2: Example Partition Configuration for a Linux File Server

Filesystem	Size (MB)	Mounted Directory
/dev/sda1	100	/boot
/dev/sda2	6000	/
/dev/sda5	4000	/var
/dev/sda6	8000	/usr
/dev/sda7	2000	Swap space
/dev/sda8	20000	/home
/dev/sda9	6000	/home/shared

The /usr filesystem is large enough to include key services such as Samba and the Linux graphical user interface. Most of the disk space has been allocated to /var for the log files and for FTP and Web services, to /home for individual user files, and to /home/shared for common files. If there's room left over, you can configure these directories on logical volumes and add to them as your needs evolve. Of course, this is only an example. The amount of disk space you allocate for file sharing will depend on factors such as the number of users and the type of files they use.

Linux Swap Space

Linux uses the swap space configured on one or more hard drive partitions to store infrequently used programs and data. Swap space can extend the amount of effective RAM on your system. However, if you don't have enough actual RAM, Linux may use the swap space on your hard drive as virtual memory for currently running programs. Because hard drive access can be 1/1,000,000th the speed of RAM, this can cause significant performance problems.

On the Job

The relative speeds of RAM and hard drives are evolving; in many cases, hard drive access times are fast enough that large amounts of swap space have lower performance penalties. However, the rule of thumb still applies: RAM is much faster than hard drives.

But you can't just buy extra RAM and eliminate swap space. Linux moves infrequently used programs and data to swap space even if you have gigabytes of RAM.

Normally, Linux (on a 32-bit Intel-style computer) can use a maximum 4GB of swap space in partitions no larger than 2GB. This 4GB can be spread over a maximum of eight partitions. The typical rule of thumb suggests that swap space should be two to three times the amount of RAM. However, at larger amounts of RAM, the amount of swap space that you need is debatable.

The way Red Hat assigns default swap space is based on the amount of RAM on your system and the space

available in your hard drive. As discussed earlier, graphical installations of RHEL require at least 192MB of RAM. If your system has the minimum amount of RAM and there's room available on your hard drives, Anaconda configures a swap partition of twice this size (384MB). For Intel 32-bit systems, Red Hat suggests a swap partition at least equal to the amount of RAM on your system. But it isn't required; I have a couple of 2GB systems for which 1GB of swap space is more than sufficient.

On the Job

Red Hat RAM and swap space requirements vary if you're installing RHEL on computers with non-Intel 32-bit CPUs.

In any case, you want to make the swap space you create as efficient as possible. Swap partitions near the front of a hard disk, thus on a primary partition, have faster access times. Swap partitions on different hard drives attached to separate disk controllers give Linux flexibility as to where to send swap data. Linux can start a program through one hard drive controller and move files to and from swap space on a separate hard drive controller simultaneously.

BIOS Limits

Some computers built before 1998 may have a BIOS that limits access to hard disks beyond the 1024th cylinder. Some older BIOSs report only 1024 cylinders on a hard drive no matter how many actual cylinders are present. Computers that are subject to this limit can't see partitions beyond this cylinder. In this case, you should configure the Linux /boot directory on its own partition. Make sure that partition is located within the first 1024 cylinders of the hard drive. Otherwise, the BIOS won't be able to find the partition with the Linux kernel.

Most PCs manufactured for a few years after 1998 have a built-in fix called *logical block addressing,* or *LBA*. A system that can report LBA will adjust the cylinder, head, and sector numbers such that the entire disk is available using these logical addresses. This has been superseded by the Enhanced BIOS, which virtualizes this form of addressing.

Multiple Controllers

It is possible and desirable to use more than one disk controller interface card at the same time on the same PC. This is a common method to increase throughput on your system by reducing your read/write bottlenecks to the only disk.

You can use both SATA/SCSI and ATA/IDE controllers in the same machine (which is the situation on my desktop system), but you should be aware of a few snags. Older BIOS may have access only to the first two IDE hard drives. Also, SCSI disks may not be accessible if IDE drives are installed. The BIOS might have a setting to allow you to boot from SCSI hard disks-or even USB keys. Make sure you understand which drives the BIOS will be able to access, because if you install /boot on an inaccessible drive, the BIOS won't be able to find your Linux boot files.

Certification Objective 2.07-Post-partition Installation Steps

Naturally, there's more to installation than partitioning. You'll need to configure networking, set the root password, set the current timezone, select basic package groups, and customize additional package groups if desired. And after installation, you have to deal with the First Boot process, but that's for later in this chapter.

If you've accepted the default and haven't selected the Review And Modify Partitioning Layout option described earlier, you won't see a boot loader configuration step and should skip to the "Networking" section a bit later in the chapter.

The Boot Loader

Next, you'll be able to configure the boot loader, as shown in Figure 2-9. This can help you configure how your BIOS finds Linux (and possibly other operating systems) on your computer. GRUB is the default, which is described in more detail in Chapter 3.

Figure 2-9: Configuring a boot loader

In this screen, you'll need to make several decisions:

1. If you already have another boot loader (such as the one associated with VCOM System Commander or Symantec Partition Magic), click No Boot Loader Will Be Installed.
2. If you want to change how an operating system appears in the GRUB menu, select it and click Edit. For example, when I'm dual-booting with Microsoft Windows XP, that boot option is shown as "Other." In that case, I click Other, click Edit, and change the Label name to Windows XP.
3. If you select Use A Boot Loader Password, it immediately opens the Enter Boot Loader Password window, where you can set up this option.
4. If you select Configure Advanced Boot Loader Options, you're taken to a different menu before the next step, where you can place the boot loader on the first sector of the boot partition and force LBA32

addressing, needed for some older hard drives.

5. You can also add kernel parameters of your choice with the Advanced Boot Loader Configuration menu.

Networking

Now you'll be able to configure this computer on your network. Assuming you've set up installation from a remote computer on the network, you'll see the settings you entered previously, as shown in Figure 2-10. You can either configure the IP address information shown manually, or you can leave this task to a DHCP server.

Figure 2-10: Configuring networking

If you want to change the characteristics of a network card (many computers have more than one), select it and click Edit. This opens the Edit Interface window associated with the network card. You can even enable IPv4 and/or IPv6 address support. Some DHCP servers can assign host names as well. Make your selections and click Next to continue.

Exam Watch

For the RHCE and RHCT installation exams, follow the IP address instructions carefully. It's possible that you may be told to leave configuration to a DHCP server and set up a specific host name; alternatively, you might set a static IP address for the network gateway and DNS servers.

Time and Root Passwords

Setting a timezone is more important than just making sure you have the right time. Web sites with servers in different timezones may have to synchronize clocks. Accuracy in e-mail timestamps can help your users keep their correspondence straight.

There are two questions associated with the timezone section of the installation process. First, you need to identify what timezone you're in. Even though there are only 24 hours in a day, there are many more than 24 different timezones. Some areas move to and from daylight savings time on different dates, or sometimes not at all.

Unless you're dual-booting with an operating system such as Microsoft Windows, you should activate the System Clock Uses UTC option. UTC is a French acronym, which is the atomic realization of what is slightly inaccurately known as Greenwich Mean Time (or Zulu Time).

After you set the timezone, the next step is to set the root password. Don't forget it, or you'll have to use one of the techniques described in Chapter 16 to rescue your system and restore your root password.

Baseline Packages

In most cases, it's best to *modestly* customize the packages that you install during the Red Hat exams. Alternatively, you can use the pirut tool, described in more detail in Chapter 5. When you see the screen shown in Figure 2-11, you'll get to customize the package list slightly. From this screen, you can choose whether to install the Software Development and Web Server package groups.

Figure 2-11: Basic package customization

But this assumes a default installation of RHEL. What you actually do install is based on two factors:

1. Whether you're installing from Server or Client installation media.
2. Your subscription key, entered near the beginning of the installation process.

In addition, the package groups mean different things depending on whether you've downloaded the Client or Server versions of RHEL. For example, while the Office and Productivity package group from the Client includes the OpenOffice.org suite, the package group of the same name from the Server includes only a couple of document readers.

We won't examine every iteration; suffice to say that some logic lies behind the subscriptions. For example, if your subscription key includes Virtualization, the default package groups installs the Xen kernel. However, surprisingly, the standard Server installation does not install standard servers such as those associated with the Web Server or FTP Server package groups.

The next two sections focus on package groups. If you see requirements on the exam for a mail server, graphics applications such as The GIMP, and recompiling the kernel, you'll want to select the Mail Server, Graphics, and Development Tools package groups.

But don't overdo it. During the GUI installation process, you can customize optional packages to be installed one by one, but the time it takes to read and choose these packages is usually not worth the time saved during the actual installation.

Exam Watch

Unless you have a specific requirement for Xen-based virtualization, don't install the Virtualization package group. As of this writing, it leads to very different settings in the boot loader; Xen-based kernels cannot yet handle all of the hardware of a standard Linux kernel.

Red Hat package groups are organized logically; for example, all the packages associated with the GNOME desktop environment belong to one Red Hat package group. It's important to choose only the package groups you need. Fewer installed packages means more room for personal files for you and your users, as well as the log files you need to monitor your system and actually get some use from your applications. On the exam, fewer installed packages leaves more time to configure the required services.

On the Job

Understanding how these package groups work is important in a kickstart installation, which is described in more detail in Chapter 5.

Package Groups

This section includes the briefest possible overview of each of the packages you can select during the RHEL installation process. Remember that installation of some of these packages depends on installation of others; for example, if you want to install the GNOME Desktop Environment package group, the Red Hat installation program will make sure that you install the X Window System package group as well. As you can see from Figure 2-12, there are high-level groups, such as Desktop Environments, and regular package groups, such as the GNOME Desktop Environment.

Figure 2-12: Red Hat Enterprise Linux package groups

For complete details of the RPMs associated with each package, go to the first RHEL installation CD and read the comps-rhel5-server-core.xml file in the /Server/repodata directory in the text editor or Web browser of your choice. If you're installing the RHEL 5 client, substitute the comps-rhel5-client-core.xml file in the /Client/repodata directory.

These packages, as well as the order in which they are presented, are based on RHEL. If you're using Fedora Core or one of the third-party rebuilds, the packages may vary. In any case, the best way to study what's in each package group is through the graphical installation.

For example, Figure 2-13 illustrates the RHEL installation, with a focus on the Mail Server package group. As you can see, sendmail is installed by default; if you need to install postfix, you'll need to make changes.

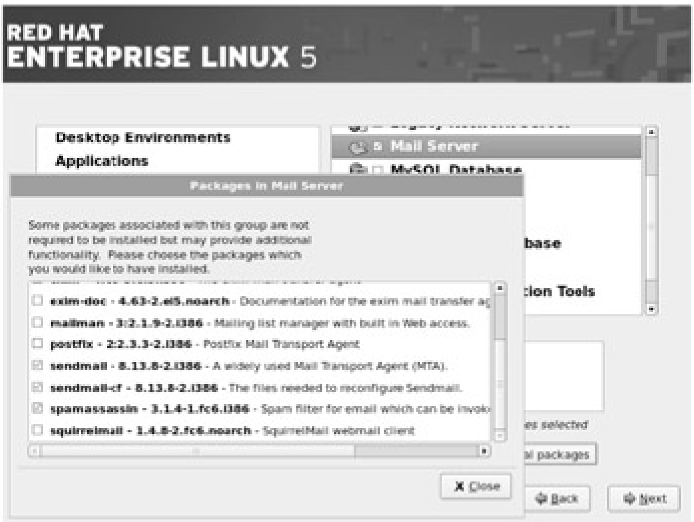

Figure 2-13: Red Hat Enterprise Linux Mail Server package group details

Take some time studying this screen. Examine the packages within each package group. You'll learn about the kinds of packages that are installed by default. If you don't add them during the installation process, it isn't the end of the world. You can still add them with the **rpm** commands or the pirut tool both described in Chapter 5. What you learn here can help you select the package groups to install during the RHCE or RHCT exam. But don't overdo it; I've seen people spend (waste?) an hour customizing every last package when they can connect to the installation source to install more after installation.

In the following sections, I describe each package group in more detail, based on what you see during the RHEL Server graphical installation process. If you're installing the RHEL Client, what you see will vary in many cases.

Desktop Environments

There are two package groups in this category, the GNOME and KDE Desktop Environments. Naturally, these install the two major desktop environments associated with Linux. And if you install one, it'll install the X Window System package group, which is described shortly as one of the Base packages.

Inside the Exam

During the Installation Exam

Even if you're taking the RHCE exam, pay attention to the software associated with the Linux Desktop Environment. The RHCE exam includes RHCT components, which means that you also need to know how to set up Linux as a client.

That also means you'll be installing a number of package groups that you would not install on a computer that's being used only as a server. When you take the exam, read the configuration requirements carefully. Don't be surprised if you see a requirement to install software such as video players and the OpenOffice.org suite.

Whatever you do, don't install everything. If you don't need to install the OpenOffice.org suite and are installing from the Client CDs/DVD, you can deselect the Office/Productivity package group and save several minutes during the Installation and Configuration portion of either exam. The time you save could allow you to configure a few more services, which could determine whether you pass.

If you make a mistake during the installation process, don't panic. You can use the pirut tool after installation to add any package groups that you missed.

GNOME Desktop Environment

The GNOME group includes the basic packages required to install the GNOME Network Object Model

Environment. While GNOME is the default GUI for RHEL, read the instructions on your exam carefully. It's possible that you'll be asked to install the other major GUI, the KDE Desktop Environment.

When you choose to install the GNOME Desktop Environment package group, all but a few GNOME packages are installed by default.

KDE Desktop Environment

The KDE group includes the basic packages required to install the K Desktop Environment, which is the main alternative GUI for RHEL. It is the default GUI for a number of other Linux distributions.

When you choose to install the KDE Desktop Environment package group, all but one of the KDE packages are installed by default.

Exam Watch

Read the instructions on the RHCE and RHCT installation exam carefully. For example, if it requires you to set up only KDE, it's a waste of time to accept the default GNOME Desktop Environment!

Applications

There are a number of package groups associated with applications. The groups range from Authoring and Publishing to Graphics to Text-based Internet.

Authoring and Publishing

The Authoring and Publishing group includes support for several documentation systems, such as LinuxDoc, DocBook, and TeX.

Editors

This group include the basic text editors associated with Linux: vi and emacs. While it's essential that you know vi to use the Linux rescue mode, the emacs text editor may be the most popular text editor in the world of Linux and Unix.

Engineering and Scientific

RHEL includes a group of packages for mathematical and scientific purposes, such as gnuplot, pvm, and units.

Games and Entertainment

Be careful with this package group. Do you really want to install games on a business computing system? Some believe that computer games are useful to help newer users become comfortable with Linux. While I doubt that you'll ever have to install this package group during the RHCE or RHCT exams, read the instructions that come with your exam.

Graphical Internet

The Graphical Internet package group includes the Firefox Web browser, the GNOME-based gFTP client, and XChat.

Graphics

This package group automatically incorporates the X Window package and a number of graphical applications. This includes the most prominent Linux graphics application, The GIMP.

Office/Productivity

If you're installing from an RHEL Server source, this group includes a PDF reader (evince) and a DVI reader (tetex-xdvi). If you're installing from an RHEL Client source, this package group includes OpenOffice.org as well as related packages.

Sound and Video

Not surprisingly, the Sound and Video group installs the packages required to allow you to use sound cards and interconnect the basic components of your sound and video system: sound card, speakers, microphone, and CD/DVD drive.

Don't dismiss this package group out of hand; I've heard that some people are asked to configure a sound card during the RHCE exam. Therefore, it's possible that you'll want to install this package group when you configure the Linux Desktop Environment.

Text-based Internet

Linux includes a number of different text-based clients for Internet access, including the **elinks** Web browser; and the fetchmail and mutt e-mail readers. This is closely related to the Graphical Internet package group.

Development

Red Hat has organized all development package groups into this category. They're not just for developers; if you need to recompile the kernel, you'll need the Development Tools package group.

Additional development tools are included when you install other packages such as GNOME Development, Graphics, Web Server, News Server, and more. Unless you need to use development tools to recompile the kernel, I don't believe that you'll have to install this package group during either exam.

Development Libraries

The Development Libraries package group includes systems that can help you add or modify features for a wide range of programs such as those associated with PCI utilities and even the **rpm** command.

Development Tools

This group includes a large number of development tools, such as make, gcc, perl, and python, useful for compiling such things as the kernel.

GNOME Software Development

The GNOME group includes the basic packages required to develop additional GTK+ and GNOME GUI applications. Some of these packages can help work with a GUI tool that can help compile the kernel. For more information, see Chapter 8.

Java Development

This group includes packages that can help you develop programs in Java.

KDE Software Development

The KDE group includes the basic packages required to develop additional QT and KDE GUI applications. Some

of these packages can help work with a GUI tool that can help compile the kernel. For more information, see Chapter 8.

Legacy Software Development

Red Hat makes it possible to develop software on RHEL for older versions of Red Hat Linux. The Legacy Software Development group includes support for older C and C++ language compilers. As there is no reference to these software development packages in any materials related to the RHCT or RHCE exams, I don't believe that you'll have to install this package group during either exam.

Ruby

The Ruby programming language is becoming more popular, especially for Web development. The associated package group adds a development environment for this language.

X Software Development

The X Software Development group includes the basic packages required to develop additional GUI applications. No reference to X Software Development packages appears in any materials related to the RHCT or RHCE exams; therefore, I don't believe that you'll have to install this package group during either exam.

Servers

Naturally, there are a number of servers associated with Linux. They're all available from the Servers group. As none are selected by default, you'll want to pay careful attention here. For example, if your exam requires you to install sendmail, bind, and Samba, you'll want to install the Mail Server, DNS Server, and Windows File Server package groups.

DNS Name Server

The DNS Name Server group includes the tools you need to configure and maintain a Domain Name System server on the local Linux computer. In the Linux world, a DNS server is also known as a nameserver, based on the Berkeley Internet Name Domain (BIND). While it is not installed by default, the Red Hat Exam Prep guide suggests that you may have to configure a DNS caching or slave server during the RHCE exam.

FTP Server

This includes the default Red Hat FTP server, the Very Secure FTP daemon (**vsftpd**). While it is not installed by default, it is possible that you'll have to configure an FTP server during the RHCE exam. It also happens to be the server that Red Hat uses for its own FTP sites.

Legacy Network Servers

Some Linux gurus discourage use of the several legacy network servers due to security concerns. Nevertheless, they remain popular. They include packages that allow you to install an RSH (Remote Shell), Telnet, and TFTP (Trivial File Transfer Protocol) server. However, it's possible that TFTP servers will become more popular in the future, as they are used for automated PXE-based installations of RHEL.

Mail Server

This group includes the packages required to configure a sendmail-based IMAP or a postfix mail server. While a mail server is not installed by default, it is possible that you'll have to configure a mail server such as sendmail or postfix during the RHCE exam.

MySQL Database

The Structured Query Language (SQL) is one of the basic database languages. This group includes support for the MySQL database system. It is not installed by default. As there is no reference to SQL or databases in any materials related to the RHCT or RHCE exams, I don't believe you'll have to install this package group during either exam.

News Server

This is a simple group, incorporating the Internet Network News (INN) server. It is not installed by default, and there is no reference to a News Server in the Red Hat Exam Prep guide.

Network Servers

This package group includes a number of smaller servers that are useful for running a network, including those associated with DHCP and Network Information Service (NIS). It is not installed by default. When you select this package group during the RHEL installation process, Red Hat installs DHCP and the NIS server.

However, if you're asked to install other components of this package group during the RHCE Installation exam, you may want to customize details during the Anaconda installation process; for example, a DHCP server is not installed unless you so specify, such as is shown in Figure 2-14. Alternatively, you can just use the **rpm** or yum command as described in Chapter 5 to install the appropriate packages after RHEL is installed. Servers with similar functionality are included in the Legacy Network Servers package group.

Figure 2-14: Network Servers package group

PostgreSQL Database

SQL is one of the basic database languages. This package group includes support for the PostgreSQL database system. It is not installed by default. As no reference is made to SQL or databases in any materials related to the RHCT or RHCE exams, I don't believe that you'll have to install this package group during either exam.

Printing Support

RHEL currently includes support for the Common Unix Printing System (CUPS). It supports the next-generation printing protocol, known as Internet Printing Protocol (IPP). Once installed, CUPS can help detect network printers, with the help of the **system-config-printer** tool.

Printing Support is the only package group installed by default with most RHEL installations.

On the Job

Red Hat no longer includes the Line Print Daemon (LPD) in its latest Linux distributions.

Server Configuration Tools

Red Hat has developed a series of GUI server configuration tools. For expert users, it's faster to configure most services from the command line interface. In fact, I encourage you to learn to configure all Linux services in this way; in the long run, you'll be a better administrator.

These tools are not installed by default in RHEL. As the RHCE exam requires you to configure servers, I encourage you to install this package group for that exam. It includes about 6MB of files, so the time penalty during the installation process is trivial. If you're less certain about your skills in one or more of these services-or if nerves affect your skills during an exam-these GUI tools can be a lifesaver.

- **system-config-bind** DNS-The Red Hat BIND DNS Configuration tool
- **system-config-boot** A graphical interface for configuring the boot loader
- **system-config-httpd** Apache configuration tool
- **system-config-nfs** NFS Server Configuration tool
- **system-config-samba** Samba Server Configuration tool
- **system-config-securitylevel** Security Level Configuration tool
- **system-config-services** Service Configuration tool
- **system-switch-mail-gnome** A GUI Interface for the Mail Agent Transfer Switcher

Web Server

The Web Server group installs Apache, Squid, and the extensive array of supporting modules and configuration files. It's quite possible that you'll have to configure at least a Web server for the RHCE exam.

Windows File Server

This group includes the Samba packages required to set up Linux as a client and as a server on a Microsoft Windows-based network. It is installed by default, and it's possible that you'll have to configure Samba during the RHCE exam. RHEL includes a stable version of the Samba 3 file server.

While I haven't seen any Microsoft Windows computers at Red Hat, it is possible to configure and test Samba clients and servers using a second Linux computer as a client.

Base System

The Base System includes standard packages associated with Linux. Not all package groups in this category are installed by default.

Administration Tools

Red Hat has developed a series of GUI administration tools. For expert users, it's faster to configure most services from the command line interface. In fact, I encourage you to learn to configure all Linux services in this way; in the long run, you'll be a better administrator.

However, these tools are installed by default in RHEL. For the purposes of the RHCE and RHCT exams, I encourage you to accept the default to install this package group. If you're less certain about your skills in one or more of these services-or if nerves affect your skills during an exam-these GUI tools can be a lifesaver:

- **authconfig-gtk** Supports configuration of NIS, LDAP, and Samba clients and more; also known as the Authentication Configuration tool. It can help you connect to these network directory services, as described in the Exam Prep guide. It can also be started with the **system-config-authentication** command.
- **pirut** Lets you install packages after RHEL is installed. Also known as the Package Management tool; it can also be started with the **system-config-packages** command.
- **sabayon** Helps maintain user profiles in the GNOME Desktop Environment.
- **setroubleshoot** Helps diagnose problems associated with SELinux.
- **system-config-date** Allows you to configure the time and date of your system; also known as the Date/Time Properties tool.
- **system-config-kdump** Supports GUI-based configuration of what happens in the event of a kernel crash dump.
- **system-config-keyboard** Lets you select a different keyboard; also known as the Keyboard Configuration tool.
- **system-config-kickstart** Opens a GUI for customizing a kickstart file; also known as the Kickstart Configurator.
- **system-config-language** Supports configuration of the GUI in different languages; also known as the Language Selection tool.
- **system-config-lvm** Supports configuration of and changes to logical volumes, which can help "add, remove, and resize logical volumes," as described in the Exam Prep guide.
- **system-config-network** Allows detailed configuration of network devices.
- **system-config-rootpassword** Allows you to change the root password.

- **system-config-soundcard** Automatically configures most sound cards.
- **system-config-users** Supports creating and modifying users and groups; also known as the Red Hat User Manager.

Closely related to Pirut is Pup, the package updater. As described in Chapter 5, Pup is used to keep systems up to date. Pup is explicitly cited in the Red Hat Exam Prep guide as a tool to keep RHEL systems up to date.

Base

The Base package group includes the fundamentals required for Linux. Naturally, it's a large group with more than 100 packages. If you're installing Linux, don't deselect this group.

Dialup Networking Support

The Dialup Networking Support package group includes the packages required to support connections over telephone and ISDN modems, as well as other PPP connections. Some DSL connections also require PPP.

Java

Now that Sun Microsystems has released Java under the GPL, this programming language is set to become more important in Linux.

Legacy Software Support

The Legacy Software Support package group is not installed by default. If you install systems that require the support of older programming libraries, you may find them here.

System Tools

This package group includes a varied array of tools, from the ethereal network traffic reader to the zsh shell. This package group is not installed by default, and it's unlikely that you'll have to install this package group during either exam. While the Samba client is part of this package group, you can also install it through the Windows File Server package group. If you need to create a local repository, you'll need the createrepo RPM package.

While this package group is not installed by default, it's a good idea to review the details of this package group during your studies, just in case you need one of these packages during your exam.

X Window System

This package group includes a number of basic Linux GUI fonts, libraries, and critical tools such as the Red Hat Display Settings tool, which you can start using the **system-config-display** command. It's required if you install the GNOME or KDE Desktop Environments.

As a Linux administrator, you may have confidence in your ability to configure Linux from the command line. In practice, you may install Linux on a number of computers without the GUI. However, even the RHCE exam includes an RHCT component, which tests your ability to "Configure the X Window system and a Desktop Environment," which by definition includes a GUI.

Exam Watch

Since the RHCT exam requires you to configure a workstation, and you have to meet all RHCT requirements during

the RHCE exam, expect to install the GUI.

Exam Watch

While installation proceeds, you'll have a bit of "dead time." You can use this time to start configuring your RHEL system. Just press the CTRL-ALT-F2 command, and you'll see a shell. You'll find the standard root directory (/) mounted on the /mnt/sysimage subdirectory during the installation process. You can edit the files of your choice as soon as they're installed.

Certification Objective 2.08-Post-installation, Security, and the First Boot Process

Now you've selected the package groups you need. The installation program takes a moment (or even a few minutes) to collect the list of packages and installs any other packages on which they depend (also known as dependencies). Once complete, Linux reboots. The first time RHEL boots, it starts the First Boot process shown in Figure 2-15, which takes you through the license agreement, firewall configuration, SELinux activation, kdump configuration for kernel crashes, the date and time, software updates from the Red Hat Network, creating a user, configuring a sound card, and installing from additional CDs if required. You may not see all of the options shown in the figure; for example, if your computer does not have a sound card, you won't see that option.

Figure 2-15: First Boot configuration

If you haven't installed the GUI, what you see is quite different, as shown in Figure 2-16. We'll describe the text-mode First Boot process at the end of this section.

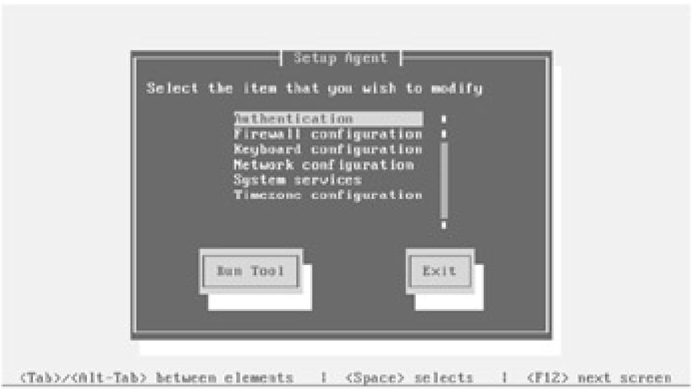

Figure 2-16: Text-mode First Boot configuration

Licensing

If you're running RHEL, you'll need to agree to the license agreement. If you do not agree, you're prompted to shut down your system and uninstall the operating system.

Initial Firewall Configuration

Next, you'll be able to configure a standard firewall for your computer. Generally, you won't need to configure a firewall for a workstation inside a LAN. Firewalls are generally located on computers that serve as junctions, or

routers between networks such as a LAN and the Internet. Figure 2-17 illustrates a configuration with two network cards that are presumably connected to different networks. The options you see here are identical to those shown in the **system-config-securitylevel** (Security Level Configuration) tool, which is described in Chapter 15.

Figure 2-17: Configuring a firewall

Initial SELinux Configuration

After configuring a firewall, you can set up basic SELinux protections on your system. It provides a different layer of defense to protect a wide variety of systems on your computer. If you don't know SELinux well, you may want to disable its protection temporarily. It's easy to do here, as shown in Figure 2-18; you can change it and customize it further using the new SELinux Management Tool described in Chapter 15.

Figure 2-18: Configuring SELinux

kdump

The kdump service allows you to configure what happens in the event of a kernel crash. In the associated First Boot screen, you can dedicate a specific amount of RAM to the process. Be aware that any such RAM is then unavailable for other processes.

Date and Time

In the Date and Time screen, you can set the date and time for your system. Under the Network Time Protocol tab, if you select the Enable Network Time Protocol option, you can synchronize your computer with a Network Time Protocol server. Red Hat provides three: 0.rhel.pool.ntp.org, 1.rhel.pool.ntp.org, and 2.rhel.pool.ntp.org. If you're not sure, you can return to this configuration screen with the **system-config-time** utility. Make any selections required by your exam, and click Next.

Exam Watch

The Red Hat exams are closed book. While you can use available documentation such as the man pages, don't expect to have a connection to the Internet during your exam.

Exam Watch

Red Hat has recently included NTP in the list of services associated with the RHCE exam, so it's quite possible that you'll have to configure a connection to a time server. For more information, see the Date/Time Configuration tool as described in Chapter 13.

Set Up Software Updates

If you want to register with the Red Hat Network, select Yes, I'd Like To Register Now, and you'll first see the Choose Server screen. You can select whether you receive updates directly from the Red Hat Network or from a local Red Hat caching service, such as a Red Hat Network Proxy or Satellite Server. If it's the latter, you'll need the URL of the local server.

On the Job

Naturally, if you're using a rebuild distribution, the information on this screen will be different-or may not even exist.

If you click the Advanced Network Configuration button, you can set your system to read through any installed local proxy server.

Then you'll see a Red Hat Network configuration screen where you can set up a connection. You'll either need the username and password of your Red Hat Network account or you'll want to click Create A New Account.

The First Boot process then collects a profile of your hardware and currently installed packages and assumes that the current IP address is the name of your system. You can review what it will send to the Red Hat Network, change the system name, and disable transmission of the hardware or software profiles. If you do not want to register at this time, you can set up a Red Hat Network connection using the **rhn_register** command (or in the GNOME Desktop Environment, select Applications | System Tools | Software Updater, which registers your system before checking for available updates).

The First Regular User

Generally, you should configure at least one regular user account on every Linux system. Using the root account for everything is considered dangerous. In the User Account screen, you can configure a regular account. You may be required to create regular users with a specific password during your exam. You can create additional users with the system-config-users (Red Hat User Manager) tool described in Chapter 6. Create a user if required. Don't click Next yet; first review the next section.

Exam Watch

If you're required to set up users over a network, pay attention to the requirements of your exam. Click the Use Network Login button and study the Authentication Configuration window. Are these users on a central NIS server or a Samba-based Primary Domain Controller (PDC)? For more information on the Authentication Configuration window, read Chapter 6.

Password Security

First, when you create a password, Red Hat has some protections. The First Boot process won't accept a password of less than six characters. In real life, it's best to use a password with a mixture of numbers and upper- and lowercase letters, and even punctuation. I like to create passwords from a favorite phrase or sentence. For example, a user could use **Ieic3teT.** because he told you "**I** **e**at **i**ce **c**ream **3** **t**imes **e**very **T**hursday." (The period at the end is part of the password.)

In First Boot's Create User screen, there's a Use Network Login button. This opens the same window and tabs associated with the Authentication Configuration tool described in Chapter 6. Briefly, you can use it to configure connections to a NIS, LDAP, or Samba user database, encryption support such as that associated with Kerberos, Hesiod, or MD5 support, and more.

Sound Card Configuration

Normally, Red Hat automatically detects sound cards. If successful, you can click the Play Test Sound button to confirm. If more than one sound card is installed, you may see them in different tabs; you can click the tab for the other sound card and repeat the process. Click Next to continue.

Additional CDs

Finally, if you have more software to install, such as software from a Red Hat supplementary CD, you'll get to do so in the Additional CDs window. Don't do this unless required by your particular exam. Click Next to continue.

You'll now see the Finish Setup screen, which tells you that "Your system is now set up and ready to use." Click Next to finish the process.

Congratulations! Installation is now complete. As the basic installation process for Rebuilds such as CentOS (and even Fedora Core 6) is almost identical to RHEL, you can use that freely available operating system to practice for the exam. In addition to the trademarks, the only major substantive difference is the lack of access to the Red Hat Network (RHN). And there is no mention of the RHN in the Exam Prep guide.

If You Haven't Configured the GUI

The text-mode First Boot process does not require a lot. As shown back in Figure 2-16, it starts the Setup Agent, which supports access to the following text-based configuration tools:

1. Authentication
2. Firewall Configuration (no access to SELinux configuration)
3. Keyboard Configuration
4. Network Configuration
5. System Services
6. Timezone Configuration

While I encourage you to learn to configure Linux from the command line, these tools are not as capable as their

GUI-based cousins. (Of course, if you want the most capable tools, learn to edit appropriate configuration files directly from the command line.) You can learn more about the GUI versions of these tools throughout this book.

Caveat Emptor on Installation

Do not worry if you make a mistake the first time you practice installing Linux on a test computer. Just redo the installation; it will be significantly faster and easier than trying to correct a problem. With so many installation options and possibilities available, it is not possible to name them all or take them all into account here. In most cases, the default is sufficient if you do not understand the question posed. Move on and get it installed, and then read the FAQs, HOWTOs, and other related documents once you are up and running. You can always reinstall. The second and third installs are actually a good thing, considering you need to know this process very well for the Red Hat exam.

Exam Watch

If you have to reinstall Linux on the Installation and Configuration part of the exam, you may not have time to configure services as required.

Exam Watch

You are not allowed to reinstall Linux during the RHCE or RHCT Troubleshooting and System Maintenance portion of those exams

Certification Objective 2.09-Installation Validation

RHEL creates a number of files during the installation process. These files essentially document what happened. The basic installation log file, /root/install.log, lists the packages that Anaconda installed on your system. The commands used by Anaconda to install Linux are stored in the /root/anaconda-ks.cfg file. This can serve as a template for the kickstart process, which you can use to install RHEL automatically on different computers. I describe this process in more detail in Chapter 5.

The Installation Log File

The installation log file, /root/install.log, provides a baseline. After you run Linux for some time, you'll probably have installed and upgraded a number of additional packages. You can refer back to this file to find the packages installed when Linux was installed on this computer.

Installation Troubleshooting

Installation involves many running processes and many parts. The system logs everything to an installation log file and separates related information among four of the five virtual console screens supported during the installation.

Exam Watch

If your installation is trouble-free, you'll have a few minutes on your hands during the Installation part of the exam. Use that time to plan how you'll configure the services per the requirements of your particular exam. But pay attention to the following sections. If your installation gets stuck, the console screens described can quickly help you diagnose the problem.

The Console Installation Output Screens

Six consoles are available during the installation process, and each tells a different story. What you see depends on whether you install in text or graphical mode. A network graphical installation is something of a hybrid; it starts in text mode before connecting to the network source and proceeding to the graphical installation.

Text mode starts in the first virtual console. Graphical mode runs in the sixth virtual console (it used to be the seventh console). You can switch between virtual consoles using the commands defined in Table 2-3. If you're in text mode, you don't need to use the CTRL key (but it does no harm). As you can see in the table, each console is associated with a function key.

Table 2-3: Installation Virtual Console Commands and Functions

Command	Console and Function
CTRL-ALT-F1	Text installation display; if you're running in graphical mode, it includes the basic commands to start graphics drivers.
CTRL-ALT-F2	Accesses a bash shell prompt; available after the first few installation steps.

CTRL-ALT-F3	Lists the log of installation messages; if network problems occur, you may see related messages here.
CTRL-ALT-F4	Displays all kernel messages, including detected hardware and drivers.
CTRL-ALT-F5	Installation displays partition formatting; nothing is shown here until Anaconda formats the actual partitions.
CTRL-ALT-F6	Graphical installation display; active only if you're running the installation program in graphical mode (was formerly available from CTRL-ALT-F7). Naturally, if you're installing in text mode, nothing is shown in this console.

The messages on the third and fourth consoles can scroll by quickly; fortunately they're collected in dedicated files, which are described shortly.

Installation Bash

You can find a bash shell on the second console, which can help you review what has been installed so far. Check it out for yourself with the CTRL-ALT-F2 command. You'll see the following installation boot prompt during the installation process:

```
s
s h3 .1#
```

Certification Summary

One of the two parts of the Red Hat RHCE and RHCT certification exams tests your ability to install Linux in different situations. In this chapter, you learned to install RHEL over a network. You also worked with the major configuration tools that are part of the installation process.

Linux works well on most current computer hardware, and RHEL is no exception. Plug and play, ACPI, and APM systems are integrated into Linux. If you want support from Red Hat, use hardware that they've tested and certified.

RHEL 5 requires a minimum of 192MB of RAM for a graphical installation, and 512MB of hard drive space. That does not include the space required for a swap partition, user files, and more. It's possible to create a functional system (for a simple server) in a 2GB hard drive. If you install everything associated with RHEL 5, including support for various languages, you could need 10GB or more.

Linux represents hardware with devices, whether they're attached by serial or parallel ports, or hotswapped through USB or IEEE 1394 devices. In most cases, Linux can automatically mount devices such as USB keys, IEEE 1394 drives, and even digital media cards once installed.

You'll usually install Red Hat Enterprise Linux over a network. I've shown you how to set up a network installation server in this chapter. The same basic lessons apply if you're studying for the RHCE or RHCT exams.

Hard drive partition planning is quite important. How you assign partitions to directories depends on the size of your hard drives, what you plan to install, and the demands on the system. Appropriately configured partitions can prevent overloads on key systems, allow for appropriate swap space, and improve security on key files.

There are a number of ways to customize your installation. The distribution is organized in package groups. Red Hat starts with baseline package groups, which are the minimum requirements for the operating system. These include default packages for a functional client or server. You can customize by adding or subtracting the package groups of your choice. The selections you make are critical during the Red Hat installation exams.

After installation comes the First Boot process, which varies depending on whether you've installed a GUI, which can help you configure a firewall, SELinux protection, date and time, the first user, password security, and sound cards. The standard First Boot process assumes a GUI. If you haven't installed a GUI, RHEL uses a Setup Agent in its place.

The Linux installation is extremely flexible. You can troubleshoot the installation process with several different consoles. Some provide useful messages: one console provides a bash shell prompt where you can inspect the current detailed status of the installation. After Linux is installed, you can find out what happened. The /var/log/dmesg file helps you figure out what hardware was detected. The /root/install.log file lists the packages that were installed.

Understanding the installation process is one of the keys to success on the RHCE exam. Find a spare computer. Practice every installation scenario that you can imagine.

Two-Minute Drill

The following are some of the key points from the certification objectives in Chapter 2.

Hardware Compatibility

? If you have a subscription to Red Hat Enterprise Linux, it's best to use hardware tested and documented by Red Hat. Alternative sources of documentation include the Hardware Compatibility List of the Linux Documentation Project.

? Linux has made excellent progress with plug and play; if conflicts occur, you may be able to diagnose them with the help of files in the /proc directory.

CPU and RAM

? Typically, swap space should be two to three times the amount of RAM. However, the amount of swap space you need is debatable when you have larger amounts of RAM.

? Depending on your requirements, Red Hat Enterprise Linux installs between approximately 512MB and 10GB of files. This does not include swap space requirements.

? When you plan space for any RHEL installation, remember to leave room for user data, additional applications, services, and a swap partition.

Hotswap Buses

? Linux plug and play works well; in most cases, when you plug a device into a serial, parallel, USB, IEEE 1394, or PC Card port, Linux detects and adds the drivers automatically. If an external disk is present, it is automatically mounted.

? Device Management includes a number of commands that can help you find detected devices and modules, including **lsusb**, **lspci**, and **lshal**.

Configuring a Network Installation

? When practicing for the exam, you may need to install over a network.

? You can create a network installation server using NFS, HTTP, or FTP.

? You can even install locally from a hard drive partition.

The First Installation Steps

? You can usually start the RHEL installation process directly from a bootable CD or USB key.

? The installation process is fairly straightforward and self-explanatory. Default package groups depend on the installation subscription number.

? When you practice installing RHEL, don't worry if you make a mistake during the process. It is usually easiest to restart the process from the beginning.

Configuring Partitions, RAID, and LVM

? Linux has a simple naming standard for disk partitions: three letters followed by a number. The first letter reflects the type of drive (*h* for PATA/IDE, *s* for SATA or SCSI). The second letter is *d* for drive. The third letter represents the relative position of the disk. The number that follows is based on the relative position of the partition on the disk.

? The first PATA/IDE drive would be hda and the next hdb, then hdc, and hdd.

? It's helpful to configure separate partitions for important data such as Web services, databases, FTP sites, and e-mail.

? Unless you use LVM, there is no easy way to resize Linux partitions. Therefore, you need to consider your partition scheme carefully.

Post-partition Installation Steps

? After configuring partitions, you'll need to set up network devices and assign a root password.

? The baseline packages associated with different subscriptions vary; you can customize them during the post-partition installation steps.

Post-installation, Security, and the First Boot Process

? After installation is complete, RHEL reboots and starts the First Boot process.

? The First Boot process takes you through licensing, configuring firewalls, enabling or disabling SELinux, configuring date and time, adding the first regular user, setting up sound cards, and more.

? If you didn't install the GUI, the First Boot process is the setup tool that is a front end to a number of other text-based tools.

Installation Validation

? If you have trouble during the installation process, a number of log files can help.

? There are several virtual consoles that can help you validate the installation.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

Hardware Compatibility

1. In what file can you find more information associated with your CPU? ?

___

Answers

1. The file most closely associated with information on the CPU is /proc/cpuinfo. Less information on the CPU is available from other files, such as /var/log/dmesg.

CPU and RAM

2. Name three different architectures on which you can install RHEL 5. ?

___

Answers

2. There are versions of RHEL 5 available for at least six different architectures (depending on interpretations): i386, x86_64 (AMD Athlon and AMD64), ia64, IBM zSeries, IBM pSeries, and IBM iSeries.

Hotswap Buses

3. What command lists all PCI devices connected to your system? ?

___

Answers

3. The **lspci** command lists all PCI devices connected to your system. If you want more information, you can run **lspci -v** or even **lspci -vv**.

Configuring a Network Installation

4. Name three network servers that you use to serve Red Hat Enterprise Linux installation files. ?

___

___

___

5. If you're unsure on how to customize a firewall, what can you run from the command line to eliminate all current firewall rules?

____

Answers

[4.](#) Three network servers that you can use to serve Red Hat Enterprise Linux installation files are FTP, HTTP, and NFS.

[5.](#) If you're unsure how to customize a firewall, the simplest way to flush out the rules is with the **iptables -F** command.

The First Installation Steps

6. If you need to check the integrity of the installation CDs in a purchased Red Hat boxed set, what command should you run from the boot: prompt when starting the installation process?

____

Answers

[6.](#) The **linux mediacheck** command at the boot: prompt adds a step to the installation process that allows you to check the integrity of installation CDs.

Configuring Partitions, RAID, and LVM

7. Which device file is associated with the fourth logical partition on the second SATA drive?

____

8. To configure a RAID 6 array, what is the minimum number of partitions you need?

____

Answers

[7.](#) The device file associated with the fourth logical partition on the second SATA drive is /dev/ sdb8. The first logical partition is /dev/sda5.

[8.](#) To configure a RAID 6 array, you need at least four RAID partitions.

Post-partition Installation Steps

9. List three names for package groups associated with servers.

____

____

____

Answers

9. The names of package groups within the Server group that install servers are DNS Name Server, FTP Server, Legacy Network Server, Mail Server, MySQL Database, Network Servers, News Server, PostgreSQL Database, Printing Support, Web Server, and Windows File Server. While there's no requirement to memorize these names, if you've installed RHEL a few times, you should be quite familiar with at least the names of a few of these groups.

Post-installation, Security, and the First Boot Process

10. You've just installed RHEL 5 and have just rebooted your system. Why might you not see the GUI First Boot process?

____

Answers

10. If you don't see the GUI First Boot process after RHEL reboots for the first time, you probably haven't installed the GUI. (It's also possible that you have a configuration problem, such as a full partition associated with /home, that prevents the GUI from booting normally.)

Installation Validation

11. If you want access to a command line during the installation process, what keys would you press?

____

12. If you suspect a networking problem freezing the installation process, what console can you access to confirm?

____

Answers

11. If you want access to a command line during the installation process, press ALT-F2 (or if you're installing via the GUI, press CTRL-ALT-F2).

12. If you want to see any error messages associated with access to a network installation server during the installation process, press ALT-F3 (or if you're installing via the GUI, press CTRL-ALT-F3).

Lab Questions

Several of these labs involve installation exercises. You should do these exercises on test machines only. The instructions in these labs delete all of the data on a system. As suggested earlier, one option is to use a virtual machine that can simulate a computer inside your operating system. An example of this is VMware, available from www.vmware.com; or Xen, which is included with RHEL 5.

Lab 1

1. You need to test Red Hat Enterprise Linux as a replacement for your current RHL 9 installed Web server. But you do not want to lose the current RHL 9 Web setup just yet. You just want to test RHEL 5 using the Web pages and CGI scripts to see if they will work. What can you do? (Note: Fresh installations from Red Hat Linux to RHEL 5 are recommended.) ?

Answers

1. Scenario 1: Buy a new disk and add it to the system. Then do a custom install to create a new installation of RHEL to partitions on the new disk, adding an entry to /boot/grub/grub.conf to provide a boot option to both versions of Linux.

   Scenario 2: No space on server. Hmm?. You've got to get creative and either find a test computer on which you can do the test install or back up everything on the main server after taking it off line. Perform a new installation of RHEL. Copy your httpd.conf configuration file and see how it works. If it fails, you can restore everything back to the way it was. Note: Test your backups first before overwriting an existing operating system.

Lab 2

2. You want to practice network installations. To do so, set up an FTP installation server on a different Linux computer using the instructions described earlier in this chapter. These instructions also work if you want to create an FTP installation server on Fedora Core. ?

   If you don't have another Linux computer, you can set up an FTP server on Microsoft Windows 2000/XP Professional/2003/Vista for this purpose.

   For the purpose of this exercise, assume that you've been asked to install a Web server, a DNS server, an FTP server, and a mail server during the RHEL installation process.

Answers

2. As described earlier in this chapter, the standard Red Hat FTP server is vsFTP; the default location for download files is the /var/ftp/pub directory. You'll want to specify a subdirectory to copy the files from the root directory of the installation CDs.

   As this is a book on RHEL, I do not describe the steps needed to create an alternative FTP server on a Microsoft Windows computer.

   To install a Web server, a DNS server, an FTP server, and a mail server during the RHEL installation process, you need to select the DNS Name Server, Web Server, FTP Server, and Mail Server package groups.

Lab 3

3. You want to practice network installations. To do so, set up an HTTP installation server on a Linux computer using the instructions described earlier in this chapter. These instructions also work if you want to create an FTP installation server on Fedora Core.

   If you don't have another Linux computer, you can set up an HTTP server on Microsoft Windows 2000/XP Professional/2003/Vista for this purpose.

   For the purpose of this exercise, assume that you've been asked to install a Samba server and a print server, and you will need to recompile the kernel.

Answers

3. As described earlier in this chapter, the standard Red Hat HTTP server is Apache. The default location for download files is the /var/www/html directory. You'll want to specify a subdirectory to copy the files from the root directory of the installation CDs.

   As this is a book on RHEL, I do not describe the steps needed to create an alternative HTTP server on a Microsoft Windows computer.

   To install a Samba server, a print server, and the packages associated with recompiling the kernel during the RHEL installation process, you need to select the Windows File Server, Printing Support, and Kernel Development package groups.

Lab 4

4. In this lab, you will distribute your filesystem over more than just one partition-as a workstation. You will need to create the partitions on a 20GB or larger PATA/IDE hard disk (see Table 2-5). If your hard drive is larger, don't use the extra space. If your system has a SATA or SCSI drive, substitute device names (e.g., sda2 for hda2) accordingly. ?

1. Create a Linux boot CD from the boot.iso image file, and then reboot the system.
2. Select manual partitioning at the appropriate step.
3. Use Disk Druid to reconfigure the partition table.
4. Delete all partitions.
5. Create the first partition with 100MB of disk space, ext3, and assign to /boot.
6. Create the next primary partition, hda2, as Linux Swap, and assign to ID 83.
7. Create a third partition with about 5500MB of disk space, ext3, and assign it to the root directory, /.
8. Create an extended partition containing all the rest of the disk space. Make it *growable*.
9. Create the first logical partition, fifth in number, with about 4GB, and assign it to /var.
10. Create two more logical partitions, hda6 and hda7. Split the remaining space between these two partitions (about 5GB each). Set it up with to a software RAID filesystem.
11. Make a RAID 1 device from the two new software RAID partitions, formatted to ext3, and assign it to /home.

    On the Job

    In the real world, you should never configure different parts of a RAID array on the same hard drive. If you do this, the failure of any single hard drive can lead to the loss of all of your data on that array. However, you may have to do so if the computer on your exam has only one physical hard drive.
12. Continue with the installation process, using your best judgment.
13. When asked to select packages, make sure that the Office/Productivity, Graphics, Graphical Internet, and Games package groups are selected.
14.

Answers

4. No special solutions are required for these labs; they're simply intended to help you practice installing Linux in a variety of different situations. The more you practice different configurations, the faster you can set up Linux during the Installation and Configuration portion of your exam.

Lab 5

5. In this lab, you will install RHEL to create a basic server. You will need to create the partitions on a 10GB or larger hard disk (see Table 2-6). If your hard drive is larger, don't use the extra space. If your system has a SATA or SCSI drive, substitute device names (e.g., sda2 for hda2) accordingly.

Table 2-6: Custom Installation as a Server, 2 GHz Pentium, 10GB Single Disk, 256MB RAM

Partition	**Size**	**Use**	**Comment**
hda1	100MB	/boot	Maintains boot files
hda2	500MB	swap	Probably plenty of space
hda3	5GB	/	The root directory
hda4	4500MB	Extended partition	Solely a container for logical partitions
hda5	500MB	/var	For print spool files
hda6	1000MB	/var/www	Web services
hda7	2000MB	/home	No interactive users
hda8	1000MB	/usr	Additional network services

    1. Create a Linux installation USB from the diskboot.img image file or a boot CD from the boot.iso file, and then reboot the system.
    2. Make sure to boot from the new media.
    3. Select custom partitioning at the appropriate time.
    4. Delete all partitions.
    5. Create the first partition with 100MB of disk space, formatted to ext3, and assign it to /boot.
    6. Create the next primary partition, hda2, with about 500MB of disk space, as Linux Swap.
    7. Create the third partition with about 5GB disk space, Linux Native, and assign to the root directory, /.
    8. Create an extended partition containing all the rest of the disk space, 4500MB.
    9.

Answers

5. No special solutions are required for these labs; they're simply intended to help you practice installing Linux in a variety of different situations. The more you practice different configurations, the faster you can set up Linux during the Installation and Configuration portion of your exam.

Lab 6

6. In this exercise, you will install RHEL to configure the partitions for an imaginary database server. You will need to create the partitions on a 25GB or larger hard disk (see Table 2-7). The main use for such a system is as a database, file, and print server, with few interactive users. If your hard drive is larger, don't use the extra space. If your system has a SATA or SCSI drive, substitute device names (e.g., sda2 for hda2) accordingly.

    1. Create a Linux installation USB from the diskboot.img image file (assuming you can boot from the USB key) or a boot CD from the boot.iso file, and then reboot the system.
    2. Make sure to boot from the new media.
    3. Select custom partitioning at the appropriate time.
    4. When prompted, select Disk Druid to edit partitions.
    5. Delete all partitions.
    6. Create the first partition with 100MB of disk space, formatted to ext3, and assign it to /boot.
    7. Create the next primary partition, hda2, with about 1000MB of disk space, as Linux Swap.
    8. Create the third partition with about 10GB disk space, Linux Native, and assign it to / (root).
    9. Create an extended partition containing all the rest of the disk space, about 14GB.
    10. Create the first logical partition, hda5, with about 3GB, formatted to ext3, and assign it to /var.
    11. Create the next two logical partitions, hda6 and hda7, with about 3.5GB each. Format each to the software RAID filesystem.

        On the Job

        In the real world, you should never configure different parts of a RAID array on the same hard drive. If you do this, the failure of any single hard drive can lead to the loss of all of your data on that array. However, it may be necessary to do so if the test computer you're using has only one physical hard drive.
    12. Use the Make RAID option to set up a RAID 1 array from these two partitions. Format it to ext3 and assign it to /opt.
    13. Create the next two logical partitions, hda8 and hda9, with about 2GB each. Format each to the software RAID filesystem.

Answers

6. No special solutions are required for these labs; they're simply intended to help you practice installing Linux in a variety of different situations. The more you practice different configurations, the faster you can set up Linux during the Installation and Configuration portion of your exam.

PREV NEXT

Chapter 3: The Boot Process

When you've finished reading this chapter, you'll know the fundamentals of the boot process. When RHEL 5 is properly installed, the BIOS points to the GRUB boot loader, normally on the appropriate master boot record (MBR). GRUB points to and initializes the Linux kernel, which then starts init, the first Linux process. The init process then initializes the system and moves into appropriate runlevels. When Linux boots into a specific runlevel, it starts a series of services. You can customize this process.

Certification Objective 3.01-The BIOS Initialization Sequence

While not officially a Red Hat exam prerequisite or requirement, a basic understanding of the BIOS is a fundamental skill for all serious computer users. While many modern computers allow you to boot directly from the media of your choice, such as an RHEL 5 installation CD or a rescue USB key, that may not be possible during your Red Hat exam. Therefore, you need to know how to modify the BIOS menu to boot from the media of your choice.

Inside the Exam

Understanding the Boot Process

Both Red Hat (RHCT and RHCE) exams require intimate knowledge of the boot process. If you have problems booting into the default GUI, you need to know how to boot into a different runlevel-and what you can do in key runlevels. From the Red Hat Exam Prep guide, the associated RHCT skill is

- Boot systems into different runlevels for troubleshooting and maintenance.

When you know GRUB, you can use it to boot RHEL into the runlevel of your choice, which can boot the system into a configuration in which you can address other issues described in this book. When you know how to modify the GRUB boot loader, you can change the way Linux boots on your system.

From the Red Hat Exam Prep guide, the associated RHCE skill is

- Diagnose and correct boot failures arising from bootloader, module, and filesystem errors.

The focus in this chapter is the boot process, and therefore the boot loader. Problems associated with kernel modules and filesystems are addressed in Chapters 4, 8, and 16.

Because of the variety of BIOS software available, this discussion is general. It's not possible to provide any sort of step-by-step instructions for modifying the wide array of available BIOS menus.

Basics of the BIOS

When you power up a computer successfully, the first thing that starts is the BIOS. Based on settings stored in stable, read-only memory, BIOS performs a series of diagnostics to detect and connect the CPU and key controllers. This is known as the Power On Self Test (POST). If you hear beeps during this process, you may have a hardware problem such as an improperly connected hard drive controller. The BIOS then looks for attached devices such as the

graphics card. After the graphics hardware is detected, you may see a screen similar to Figure 3-1, which displays other hardware as detected, tested, and verified.

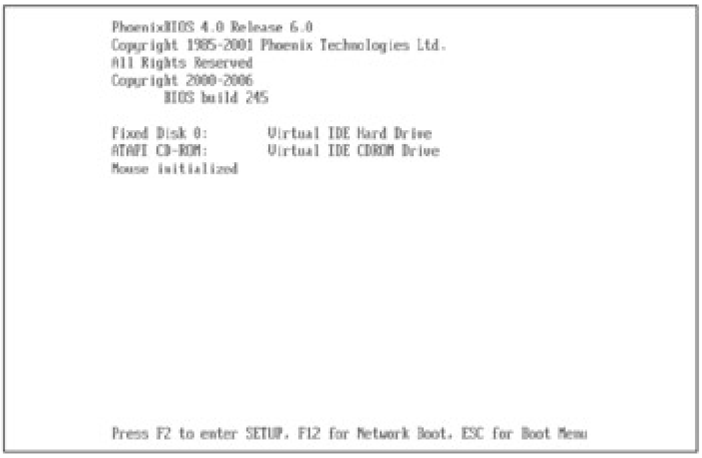

Figure 3-1: The BIOS initialization process

Once complete, the BIOS passes control to the MBR of the boot device, normally the first hard drive. At this point, you should see a boot loader screen.

Using the BIOS Menu

Generally, the only reason to go into the BIOS menu during the Red Hat exams is to boot from different media, such as a CD, floppy, or USB key. In many cases, you can bypass this process. Return to Figure 3-1. The options for this particular system are shown in the bottom of the screen. In this case, pressing F2 enters SETUP, the BIOS menu; pressing F12 boots directly from a network device; and pressing ESC starts a boot menu. The actual keys may vary on your system.

In many cases, all you see after POST is a blank screen. The BIOS is often configured in this way. In that case, you'll need to do some guessing based on your experience on how to reveal the screen shown in Figure 3-1 and access the boot or BIOS menu.

If you're fortunate, the computer you use on your exam has a boot menu accessible by pressing a key such as ESC, DEL, F2, or F12, with entries similar to:

```
      Boot Menu
1
1. Removable Devices
 2
 2. Hard Drive
3
3. CD-ROM Drive
```

Certification Objective 3.02-The GRUB Boot Loader

The standard boot loader associated with Red Hat Enterprise Linux (RHEL) is GRUB, the GRand Unified Boot loader. LILO, the Linux Loader, is no longer supported. As suggested by the Red Hat exam requirements, for the RHCT exam, you need to know how to use the GRUB menu to boot into different runlevels, and diagnose and correct boot failures arising from boot loader errors.

GRUB, the GRand Unified Bootloader

Red Hat has implemented GRUB as the only boot loader for its Linux distributions. When you start your computer, your BIOS looks for the /boot directory and finds the GRUB menu, which will look similar to Figure 3-2. If you've configured your computer with multiple operating systems, you can use the GRUB menu to boot any operating system detected during the Linux installation process.

Figure 3-2: The GRand Unified Bootloader (GRUB)

If you need to do something special with GRUB, you can edit the commands. If GRUB is password protected, you'll need to start with the **p** command. Use the **e** command to temporarily edit the file. You'll see a number of basic commands that you can use to modify GRUB, as shown in Figure 3-3 and Table 3-1. You can use these commands to test different GRUB configurations. They can help you troubleshoot problems with the GRUB configuration file, without booting, editing, and rebooting your system. That could possibly help you save time during the Troubleshooting exam. But once you find the solution, make sure to record the change in the GRUB configuration file, /boot/grub/grub.conf.

Table 3-1: GRUB Editing Commands

Command	**Description**
b	Boot the currently listed operating system
d	Delete the current line
e	Edit the current line
o	Create an empty line underneath the current line

O	Create an empty line above the current line

Figure 3-3: Details of GRUB

On the Job

Other Linux distributions store the GRUB menu in /boot/grub/menu.lst. Red Hat links that file to its GRUB menu file, /boot/grub/grub.conf.

Exam Watch

If you're troubleshooting GRUB, the GRUB menu can help you experiment with changes quickly. However, if you find the solution in this way, you'll still need to change (and test) the GRUB configuration file, /boot/grub/grub.conf.

Exam Watch

If you find a problem with GRUB during the Troubleshooting and System Maintenance part of your exam and are unsure about the solution, testing changes through the GRUB menu could save you time. However, until you record the change in the GRUB configuration file (which requires that you boot Linux into some runlevel), you won't get credit for the work that you've done.

You can also use GRUB to boot other operating systems, including various versions of Microsoft Windows, as shown in Figure 3-3.

On the Job

The /etc/grub.conf file is linked to the actual GRUB configuration file, /boot/grub/grub.conf. If you edit /etc/grub.conf, changes are automatically reflected in /boot/grub/grub.conf.

GRUB Parameters

To pass a parameter to GRUB, type the **a** command in the first GRUB menu. This allows you to append the command sent to the kernel. After typing the **a** command, you might see a command line similar to the following:

```
g
grub append> ro root=LABEL=/ rhgb quiet
```

Certification Objective 3.03-Kernel Initialization and the First Process

This section provides a basic overview of the boot process that occurs after the GRUB boot loader finds the kernel. Understanding what happens here can help you diagnose a wide variety of boot problems, some of which you might see during a Red Hat exam.

Kernel Message Analysis

Just a few messages after you boot a kernel from the GRUB configuration menu, Linux hands over boot responsibilities to the kernel. If you've disabled the **quiet** directive in the GRUB configuration file, you can watch as the messages pass quickly through the screen. To review these messages, open /var/log/dmesg or run the dmesg command.

What you see depends on the hardware and configuration of your computer. Key messages include:

- The version of the kernel
- Amount of recognized RAM (which does not necessarily match the actual amount of installed RAM)
- CPUs (labeled as CPU 0, CPU 1, and so on)
- SELinux status, if active
- Kernel command line, specifying the logical volume or root filesystem label
- Freeing of memory associated with the initial RAM disk (initramfs)
- Hard drives and partitions (as defined by their device file names, such as /dev/sda or /dev/hda1)
- Network cards, as defined by their device names, such as eth0
- Active filesystems
- Swap partitions

This file is filled with potential clues. If you've booted from the wrong kernel, you'll see it here. If Linux isn't using a partition that you've configured, you'll also see it here (indirectly). If you don't see an active network card and don't see it in the dmesg file, it may signal either a missing driver or a hardware problem with your computer.

Exam Watch

Remember that the Red Hat exams are not hardware exams. If you identify a problem with a key hardware component, such as a network card (which cannot be solved by some Linux command), inform your instructor/ exam proctor.

Driver Loading

While it's possible to create and install a "monolithic" kernel without drivers, that's not the way it's done on RHEL. Most kernels on modern PCs are modular. That means when you boot, kernels are loaded, followed by drivers. While many drivers are associated with hardware components, others include key software modules, such as the ext3 filesystem. You can review loaded modules using the **lsmod** command.

Certification Objective 3.04-The First Process and /etc/inittab

Once kernels and drivers are loaded, Linux starts loading the rest of the system. This all starts with the First Process, known as init. It loads based on the parameters defined in /etc/inittab, which specifies runlevels, the system initialization script, virtual consoles, and more.

The First Process

The Linux kernel continues the boot process by calling init. The init process in turn runs /etc/rc.d/rc.sysinit, which performs a number of tasks, including network configuration, SELinux status, keyboard maps, system clock, partition mounts, and host names. It also loads the modules described in the previous section. It does even more: the default version of this file contains more than 500 lines.

/etc/inittab

The init process then determines which runlevel it should be in by looking at the **initdefault** directive in /etc/inittab. A runlevel is defined as a group of activities. For example, the entry

```
id:5:initdefault:
```

Certification Objective 3.05-Runlevels

Linux services are organized by runlevel. Some runlevels can reboot and halt Linux. Other runlevels can boot Linux with or without networking. The default runlevel for RHEL 5 boots Linux into the GUI.

Runlevels are controlled by scripts, organized in runlevel-based directories. While the default runlevel is defined in /etc/inittab, you can override the default during the boot process from the GRUB menu.

Functionality of Each Runlevel

RHEL has six basic runlevels, as defined in /etc/inittab. Each runlevel is associated with a level of functionality. For example, in single-user mode, also known as runlevel 1, only one user is allowed to connect to that Linux system. X11 mode, also known as runlevel 5, starts Linux into a GUI login screen. The Red Hat definitions for System V init runlevels are shown in Table 3-2.

Table 3-2: Red Hat Runlevels

Runlevel	**Description**
0	Halt
1	Single-user mode, for maintenance (backups/restores) and repairs
2	Multiuser, with some network services
3	Multiuser, with networking
4	Unused
5	X11, defaults to a GUI login screen; logins bring the user to a GUI desktop, with networking
6	Reboot (never set **initdefault** in /etc/inittab to this value)

Making each runlevel work is the province of a substantial number of scripts. Each script can start or stop fundamental Linux processes such as printing (**cupsd**), scheduling (**crond**), Apache (**httpd**), Samba (**smbd**), and more. The starting and stopping of the right scripts becomes part of the boot process.

It should go without saying that if you set your **initdefault** to 0, your system will shut down when Linux tries to boot. Likewise, if you set the **initdefault** to **6**, Linux will enter a continuous reboot cycle.

The default runlevel when you boot RHEL 5 is 5, GUI with networking. If you don't install the GUI on RHEL 5, the default is 3. If you have problems with one runlevel, consider booting into another. For example, if the GUI isn't working, consider booting into runlevel 1, 2, or 3. It allows you to boot into Linux, examine appropriate logs, and address related problems.

Different runlevels can be customized. For example, if you boot into runlevel 1 and have an appropriate network connection, all you need to do is run a command such as **dhclient eth0** to activate that network connection.

Exam Watch

To practice for the Troubleshooting and System Maintenance part of each exam, it can be useful to back up and then modify critical configuration files such as /etc/inittab. But remember to do this on a test computer; if you can't solve the problem, you may lose the data on that computer. Before you proceed, learn the rescue mode techniques described in Chapter 16.

Runlevel Scripts

There are a series of scripts associated with each runlevel. For example, the default runlevel is 5, and the scripts associated with this runlevel can be found in the /etc/rc.d/rc5.d directory.

Naturally, the scripts associated with other runlevels can be found in other /etc/rc.d directories:

- rc0.d
- rc1.d
- rc2.d
- rc3.d

Certification Objective 3.06-Controlling Services

Whenever you install a service on the Red Hat exams, you'll generally want to make sure that they're active when the person grading your exam boots your system. There are three basic tools used to control services: text commands, text-based tools, and the Red Hat GUI Service Configuration tool.

Service Control from the Command Line

It's generally fastest to control services at the command line. The chkconfig command gives you a simple way to maintain different runlevels within the /etc/rc.d directory structure. With chkconfig, you can add, remove, and change services; list startup information; and check the state of a particular service. For example, you can check the runlevels where the sendmail service is set to start with the following command:

```
#
# chkconfig --list sendmail
```

Certification Objective 3.07-System Configuration Files

Red Hat sets up a number of key configuration files in the /etc/sysconfig directory. You can configure them with a text editor, with text commands, or in many cases with a Red Hat graphical tool. Many of the non-network system configuration files are discussed in this section.

It's fastest if you know how to configure these systems directly using text commands or by directly editing the key configuration file. However, if you forget how to manage one or two configuration commands or files, the Red Hat graphical tools can be a lifesaver.

I address only those systems not already covered in other chapters.

Non-network /etc/sysconfig Files

Let's return to the /etc/sysconfig directory and discuss some of the non-network configuration files listed in Table 3-3. This section covers only those files that are less likely to be of interest on the Red Hat exams. More important files are covered in other chapters. Some files can be edited directly; others can be configured with other Red Hat tools discussed in the following section.

Table 3-3: Key Non-network /etc/sysconfig Files

File in the /etc/sysconfig Directory	Description
clock	Contains defaults for the system clock, including time zone, UTC, and ARC (Alpha CPU-based) settings. If **UTC=true**, the BIOS is set to the atomic realization of Greenwich Mean Time.
firstboot	If **RUN_FIRSTBOOT=YES**, then you can start the First Boot process with the firstboot command, if you're in runlevel 5.
grub	Lists the hard disk with your /boot drive, assuming you're using the GRUB boot loader.
hwconf	Lists peripherals detected by **kudzu**. Do *not* edit this file!
i18n	Sets the default language.
init	Specifies the graphics and associated colors during the boot process.
iptables	Includes the iptables firewall commands that run when you boot Linux.
iptables-config	Adds configuration for adding iptables rules.

irda	Controls infrared devices.
kernel	Specifies defaults when you update the kernel.
keyboard	Contains keyboard configuration data: **KEYBOARDTYPE**, usually **pc**, and **KEYTABLE**, usually **us**.
kudzu	Configures hardware detection during the boot process, as it relates to serial ports, DDC (Display Data Channel) between the monitor and video card, and PS/2 ports.
ntpd	Specifies synchronization with the hardware clock after synchronizing with a remote NTP (network time protocol) server.
pcmcia	Contains PCMCIA configuration data. If **PCMCIA=yes**, Linux loads PCMCIA modules on boot.
pm	Configures hibernation characteristics.
rhn/	Directory with Red Hat Network configuration.

GUI Configuration Utilities

It's important to know how to configure RHEL 5 by hand, because it's the most efficient way to control everything on your Linux system. It's faster on the Red Hat exams, where time is of the essence. There are a number of good GUI configuration tools available; almost all of them are "front ends" that edit text configuration files, which you could edit directly.

However, there's a lot to learn about Linux. Learning how to edit *all* key Linux configuration files can be more than some RHCT/RHCE candidates can handle. While you should learn how to edit these files by hand, you may not have time. You may get nervous during the Red Hat exams and forget details. In these cases, the Red Hat GUI tools can be a lifesaver.

On the Job

The text mode setup tool is a front end to a number of other tools you can view from the text console: authentication, firewall, keyboard, network, printer, service, time zone, and GUI display configuration. You can start the tool with the setup command. While they may be faster than GUI tools, some of these tools do not have the same capabilities as GUI tools.

Date/Time Properties

With the Date/Time Properties configuration tool, you can set the date, time, timezone, and NTP server for your system. You can start it in one of three ways in the GUI: run the **system-config-date** or **system-config-time** command, or choose System | Administration | Date and Time. This opens the Date/Time Properties window shown in Figure 3-8.

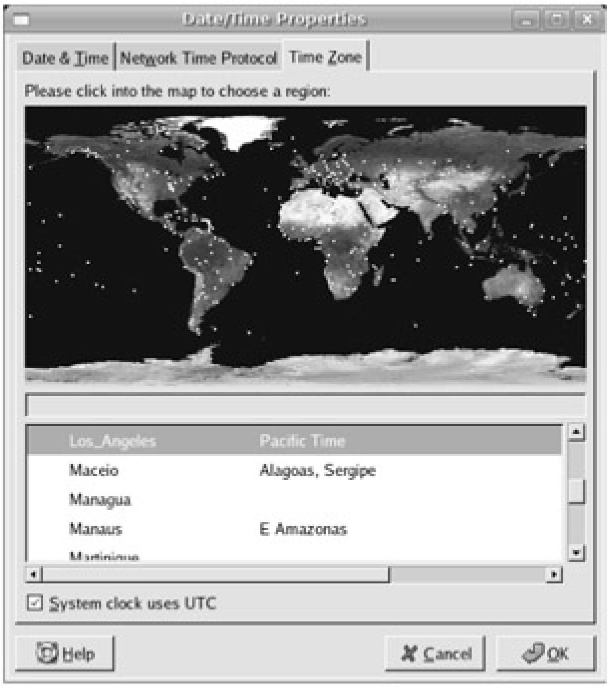

Figure 3-8: The Date/Time Properties window

Keyboard Configuration Tool

The Keyboard Configuration tool allows you to reselect the keyboard associated with your system. You can start it in one of two ways in the GUI: run the **system-config-keyboard** command or choose System | Administration | Keyboard. The options are the same as those you saw during the installation process. Results are recorded in /etc/sysconfig/keyboard.

Exam Watch

If you're studying for the RHCE, Red Hat has just added NTP configuration and troubleshooting to the Exam Prep guide. For more information, see Chapter 13.

Certification Summary

This chapter covered the basic boot process of an RHEL system. You learned the basics of the BIOS and what happens when it hands control to the GRUB boot loader. You experimented with GRUB, observing the results of various errors.

Once GRUB boots your system successfully, it hands control to the kernel. You can find out more about what happens through /var/log/dmesg and the drivers it loads. It hands control to the First Process, also known as init, as configured in /etc/inittab. It then starts various services, which you can control with various text or graphical service configuration tools.

The non-network configuration files in the /etc/sysconfig hierarchy affect basic parameters such as the system clock, kernel updates, and the keyboard. There are Red Hat GUI tools available for those who forget how to configure some key configuration files by hand during the Red Hat exams.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 3.

The BIOS Initialization Sequence

? While not strictly a part of the exam, it's important to know the basics of the BIOS.

? You can change the boot sequence from the BIOS menu.

? Once the BIOS detects your drives, it hands control to GRUB via the master boot record (MBR).

The GRUB Boot Loader

? GRUB, the GRand Unified Boot loader, is the default for RHEL 5.

? Errors in the GRUB configuration file can lead to a number of boot problems, including kernel panics.

? You can read the GRUB configuration file from the GRUB command line.

Kernel Initialization and the First Process

? You can analyze how the kernel booted your system through /var/log/dmesg.

? As the kernel initializes your system, it loads important modules such as the ext3 filesystem.

The First Process and /etc/inittab

? Once the kernel boots, it hands control to init, also known as the First Process.

? The init process starts your system as configured in /etc/inittab.

? One of the key configuration files started by the First Process is /etc/rc.sysinit.

Runlevels

? There are six different runlevels available; the default is configured in /etc/ inittab.

? Don't configure a default runlevel of 0 or 6.

?

The default runlevel configured in /etc/inittab starts scripts in the associated /etc/rc*x*.d directory, where *x* is the runlevel.

?

You can boot to the runlevel of your choice from the GRUB configuration menu.

?

Study the /etc/rc.d hierarchy and the /etc/inittab and /etc/rc.d/rc.sysinit files. This is the key to understanding what's happening during the boot process.

Controlling Services

?

The chkconfig command gives you a simple way to maintain the /etc/rc.d directory structure.

?

The **ntsysv** and **system-config-services** commands provide console and GUI tools for the same purpose.

System Configuration Files

?

There are a number of non-network configuration files in the /etc/sysconfig directory.

?

You can edit many of these files directly or use GUI tools, which you can start with commands such as **system-config-date**, **system-config-keyboard**, and **system-config-services**.

PREV NEXT

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. While the topics in this chapter are "prerequisites," it is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

The BIOS Initialization Sequence

1. If you want to boot from the CD/DVD drive, what two ways might you work with the BIOS?

____

____

Answers

1. There are two basic ways to boot directly from a CD/DVD. First, you can change the boot order within the BIOS menu. Second, in many cases, you can access a boot order menu directly with a key command such as pressing ESC or DEL, which may be shown onscreen as the computer starts up. Not all computers have a boot order menu, but if the one you're using on the Red Hat exam has such a menu, using it can save you a bit of time.

The GRUB Boot Loader

2. When you see the GRUB configuration menu, what command would you use to modify the kernel arguments? Assume a regular (non-Xen) kernel and that GRUB is not password protected.

____

3. If you've run the proper commands at the **grub>** command line, what command would you use to start booting Linux? ____________________________

Answers

2. By default, you need to press any key within 5 seconds before the default operating system is booted. When you do, you'll see the GRUB configuration menu. Press **a** to see the kernel command line; you can then modify the line adding commands. As described in this chapter, you can use this line to supersede the default runlevel.

3. The proper commands at the **grub>** command line specify the /boot directory, the kernel, the partition with the top-level root directory, and the initial RAM disk. Once executed, you can start Linux with the boot command.

Kernel Initialization and the First Process

4. What file contains the kernel initialization messages? __________________

5. What one-word command can you use to read the kernel initialization messages?

____

Answers

4. The file with the kernel initialization messages is /var/log/dmesg.

5. The one-word command that you can use to read the kernel initialization messages is dmesg.

The First Process and /etc/inittab

6. What configuration file is associated with the First Process? ?

____

7. What is normally the default runlevel for RHEL 5? ?

____

Answers

6. The configuration file associated with the First Process is /etc/inittab.

7. The default runlevel for RHEL 5 is 5.

Runlevels

8. If you can't boot into the GUI, what runlevel should you try first from the GRUB menu? ____________ ?

9. What alternatives are available from the GRUB menu to runlevel 1 if you want to boot directly into the root account? _________________ ?

Answers

8. If you can't boot into the GUI, the first runlevel you should run from the GRUB menu is 3, which is the standard command line runlevel with all standard services except the GUI.

9. The alternatives to runlevel 1 from the GRUB menu are **s** for single-user mode and **emergency** for a boot without any mounted directories or services.

Controlling Services

10. The FTP server is **vsFTPd**. If you want to make sure it starts the next time you boot into a standard runlevel, what command should you run? ?

____

11. What command would you use to open the console-based Service Configuration tool in runlevels 3 and 5? ?

____

Answers

10. If you want to make sure the **vsftpd** service starts the next time you boot into a standard runlevel, run the **chkconfig vsftpd on** command. The **chkconfig --level 35 vsftpd on** command also works.

11. If you want to open the console-based service configuration tool in runlevels 3 and 5, run the **ntsysv --level 35** command.

System Configuration Files

12. What directory contains most Red Hat system configuration files?

Answers

12. The directory that contains most Red Hat *system* configuration files is /etc/sysconfig.

PREV NEXT

Lab Questions

The Red Hat exams are unique based on their reliance on labs and hands-on demonstrations. With these questions, you're practicing the skills you need on both Red Hat exams.

Lab 1

1. In this exercise, you are going to experiment with two ways of managing services at different runlevels: the chkconfig command and the **system-config-services** utility, also known as the Service Configuration utility. The commands in this lab don't start or stop scripts immediately, just the next time you move your Linux system into runlevel 3.

    1. Open the GUI. From a text command line interface, run the **system-config-services &** command. This allows you to use the same terminal window for other commands. Alternatively, in the GNOME desktop, choose System | Administration | Server Settings | Services.

    2. The Service Configuration utility is a graphical tool for controlling the services that Linux starts and stops at each runlevel.

    3. These next steps assume that the NFS service, as controlled by the /etc/init.d/nfs script, is already running and installed. If not, pick another service to add and remove (it does not matter which as long as you restore the original condition when you're done).

    4. At the command line, run the **ls /etc/rc3.d/*nfs** command to find the priority number.

    5. Remove the NFS service from runlevel 3 using the Service Configuration utility.

    6. Switch back to the command line window and run the **chkconfig --list nfs** command to see if it has been deactivated in runlevel 3.

    7. Restore the **nfs** start script with the following command:

    ```
    # chkconfig --level 3 nfs on
    ```

Answers

1. 1.

   To open a command line interface in the default Red Hat GNOME desktop, choose Applications | Accessories | Terminal.

   2.

   In the new terminal, open the Service Configuration utility. Run the applicable graphical tool in the background so you can still use this terminal window with the following command:

```
# system-config-services &
```

Lab 2

2. In this lab, you'll move the GRUB configuration file and examine its effects on the boot process.

   1.

   Boot into Linux, and print out the contents of the GRUB configuration file.

   2.

   Move the GRUB configuration file, /boot/grub/grub.conf, to the administrative home directory. One way to do this is with the following command:

```
# mv /boot/grub/grub.conf /root
```

Answers

2. This lab should be somewhat self-explanatory. It examines what happens when a configuration file related to the boot process is missing. You'll see more examples of this in Chapter 16.

   While it's an excellent idea to experiment with other key configuration files, don't do so haphazardly. Remember to back up your configuration. And if you want to rescue your system, you may need to learn how to use **linux rescue** mode as discussed in Chapter 16.

   However, if you want to experiment with other key configuration files, back up /etc/fstab. You can change it and even move it from the /etc/ directory, in the manner described with the GRUB and /etc/inittab configuration files in this chapter. If /etc/fstab is missing, it looks like everything still works. RHEL 5 even boots into the GUI. However, the problems you'll see are subtle. No filesystem is mounted. However, you can still restore the fstab file from backup. You'll learn about mounts and /etc/fstab in more detail in Chapter 4.

Lab 3

3. In this lab, you'll see what happens without the init service. As you did in Lab 2, you'll move a key file, in this case, /etc/inittab, to the root user's home directory, and then reboot your system.

   1.

   Move your /etc/inittab configuration file to the root user home directory. You can do so with a command such as:

```
# mv /etc/inittab /root/
```

Answers

3. One of the required skills in the Red Hat Exam Prep guide is booting into different runlevels. This lab sets up a scenario in which you need to learn how to boot into a specialized runlevel, **s**, also known as single-user mode. If you want to learn more, try the **init=/bin/sh** runlevel, and see what you need to do to restore the /etc/inittab file.

Chapter 4: Linux Filesystem Administration

Overview

After installation has completed, you still have work to do. To customize the system to meet your needs, you may need to add partitions, new filesystems, automounted directories, and more. Much of this depends on the security attributes on files and directories.

While it's easier to use Disk Druid during the installation process, the fdisk and **parted** tools can help you create new partitions. Once created, filesystems can be configured on these partitions with format commands and more. Once created, you can configure mounting during the boot process in /etc/fstab or over a network using the automounter.

What you create can be protected with Security Enhanced Linux (SELinux). To that end, in this chapter you'll examine basic access control lists and set SELinux characteristics from the command line. (If you're interested in controlling SELinux using the SELinux Management Tool, see Chapter 15.)

Inside the Exam

Administrative Skills

As in the real world, it is the results that matter. It doesn't matter whether you use Disk Druid, fdisk, or **parted** to create partitions. You can create new partitions at the command line or use GUI front ends to these tools such as GParted (if it's ever included with RHEL). Make sure that your partitions meet the requirements of the exam. Just remember Disk Druid is available only during the installation process.

The current Red Hat Exam Prep guide suggests that RHCTs need to know how to *Add new partitions, filesystems, and swap to existing systems* for the Troubleshooting and System Maintenance part of their exams. It also suggests that RHCTs need to know how to

- Add and manage users, groups, quotas, and File Access Control Lists.

Remember that RHCEs also need to be prepared to do anything from the RHCT requirements.

The Exam Prep guide also suggests that RHCEs need to be prepared, during their Troubleshooting and System Maintenance sections, to

- Add, remove, and resize logical volumes.

As of this writing, SELinux has just been added to the Red Hat Exam Prep guide; in most cases, you'll need to know how to configure services to work while SELinux is running in targeted mode. However, this is a skill associated with RHCEs, not RHCTs. While the services are primarily covered in Chapters 9?15, the basic skills associated with controlling SELinux are covered near the end of this chapter.

Certification Objective 4.01-Partitioning Hard Disks

It's best to create partitions using Disk Druid during the installation process. This can save you grief as an administrator, and especially during the Red Hat exams. However, mistakes are made. You might forget to create a critical partition during the Installation and Configuration part of the exam, for example.

As suggested by the Red Hat Exam Prep guide, you may need to do more with partitions on an installed system. In the real world, you might need to create a larger /home directory partition for your users. For this purpose, the standard is still the fdisk utility, which is described shortly along with the emerging alternative, **parted**.

Before you use fdisk to create or revise partitions, you should check your free space and the partitions that are currently mounted. You can do this with the **df** and mount commands. The following example illustrates how the **df** command displays the total, used, and available free space on all currently mounted filesystems.

Note the numbers under the 1k-blocks column. In this case (except for the mounted DVD), they add up to about 35GB of allocated space. If your hard drive is larger, you may have unallocated space that you can use for another partition. Just remember to leave room for expansion in appropriate directories, such as /home, /tmp, and /var.

```
[root@Enterprise root]# df
Filesystem           1k-blocks      Used Available Use% Mounted on
/dev/mapper/VolGroup00-LogVol00
                       9903432   2333948   7058304  25% /
/dev/hda3               101105     19821     76063  21% /boot
tmpfs                   451200         0    451200   0% /dev/shm
/dev/mapper/VolGroup00-LogVol02
                       4951688    149164   4549036   4% /home
```

Certification Objective 4.02-Managing Filesystems

There are as many, if not more, filesystem types as there are operating systems. While RHEL can work with many of these formats, the default is ext3. While many users enable other filesystems such as ReiserFS, Red Hat may not support them.

Linux supports a rich variety of filesystems. Linux filesystems can be somewhat inaccurately divided into two categories: "standard" formatting and journaling. While this is an oversimplification, it suffices to describe the filesystems important to Linux. To me, a standard filesystem is an older Linux filesystem which does not log changes.

On the Job

There are a large number of filesystem types well described in the Filesystems HOWTO at www.tldp.org. Strictly speaking, there is no "standard" filesystem.

The filesystems I describe in this book are just a small list of those available for RHEL. If you have the kernel source RPMs loaded on your system, you can find a list of the filesystems supported by your kernel. For x86 systems, navigate to the /usr/src/redhat/BUILD/kernel-2.6.18/linux-2.6.18.i386 directory. Run the **make menuconfig** command and use your arrow keys to navigate to the filesystems section.

Standard Formatting Filesystems

Linux is a clone of Unix. The Linux filesystems were developed from the Unix filesystems available at the time. The first Linux operating systems used the Extended Filesystem (ext). Until the past few years, Red Hat Linux operating systems formatted their partitions by default to the Second Extended Filesystem (ext2).

There are other filesystems available for RHEL, a sample of which are included in Table 4-1. These "standard" filesystems don't include journaling features.

Table 4-1: Some Linux Standard Filesystem Types

Filesystem Type	Description
ext	The first Linux filesystem, used only on early versions of that operating system.
ext2 (Second Extended)	The foundation for ext3, the default RHEL filesystem. The ext3 filesystem is essentially ext2 with journaling.
swap	The Linux swap filesystem is associated with dedicated swap partitions. You've probably created at least one swap partition when you installed RHEL.
MS-DOS and VFAT	These filesystems allow you to read MS-DOS-formatted filesystems. MS-DOS lets you read pre-Windows 95 partitions, or regular Windows partitions within the limits of short file names. VFAT lets you read Windows 9*x* /NT/2000/XP/Vista partitions formatted to the FAT16 or FAT32 filesystems.

ISO 9660	The standard filesystem for CD-ROMs. It is also known as the High Sierra File System, or HSFS, on other Unix systems.
NTFS	The Microsoft Windows NT/2000/XP/2003 filesystem designed for username/password security. Currently supported as a read-only system.
/proc	A Linux *virtual* filesystem. Virtual means that it doesn't occupy real disk space. Instead, files are created as needed. Used to provide information on kernel configuration and device status.
/dev/pts	The Linux implementation of the Open Group's Unix98 PTY support.
NFS	The Network File System, the system most commonly used to share files and printers between Linux and Unix computers.
CIFS	The Common Internet File System (CIFS) is the successor to the Samba/Server Message Block (SMB) system based on Microsoft and IBM network protocols. Linux can use CIFS and SMB to share files and printers with Microsoft Windows operating systems.

Understanding Journaling Filesystems

As hard disks and partitions grow in size, Linux users are moving toward filesystems with journaling features. Journaling filesystems have two main advantages. First, it's faster for Linux to check during the boot process. Second, if a crash occurs, a journaling filesystem has a log (also known as a journal) that can be used to restore the metadata for the files on the relevant partition.

The default RHEL filesystem is the Third Extended Filesystem, also known as ext3. However, it isn't the only journaling filesystem option available. I list a few of the options commonly used for RHEL in Table 4-2. From this list, Red Hat officially supports only ext3.

Table 4-2: Journaling Filesystems

Filesystem Type	**Description**
ext3	The default filesystem for RHEL.
JFS	IBM's journaled filesystem, commonly used on IBM enterprise servers.

ReiserFS	The Reiser File System is resizable and supports fast journaling. It's more efficient when most of the files are very small and very large. It's based on the concept of "balanced trees." It is no longer supported by RHEL, or even by its main proponent, SUSE. For more information, see www.namesys.com.
xfs	Developed by Silicon Graphics as a journaling filesystem, it supports very large files; as of this writing, xfs files are limited to 9 × 1018 bytes. Do not confuse this filesystem with the X Font Server; both use the same acronym.

Creating Filesystems with mkfs

There are several commands that can help you create a Linux filesystem. They're all based on the mkfs command. As described in Chapter 1, the mkfs command works as a front-end to filesystem-specific commands such as **mkfs.ext2** and **mkfs.ext3**.

If you want to reformat an existing partition, take the following precautions:

- Back up any existing data on the partition.
- Unmount the partition.

There are two ways to apply formatting on a partition. For example, if you've just created a partition on /dev/sdb5, you can format it to the ext3 filesystem using one of the following commands:

```
#
# mkfs -t ext3 /dev/sdb5
#
```

Certification Objective 4.03-Filesystem Management and the Automounter

Before you can use the files in a directory, you need to mount that directory on a partition formatted to some readable filesystem. Linux normally automates this process using the /etc/fstab configuration file. When you boot Linux, specified directories are mounted on configured partitions. The mount options require some explanation, especially for removable media.

You may encounter problems if connections are lost or media is removed. When you configure a server, you could be mounting directories from a number of remote computers. You could also want temporary access to removable media such as USB keys or Zip drives. The automount daemon, also known as the automounter or autofs, can help. It can automatically mount specific directories as needed. It can unmount a directory automatically after a fixed period of time.

Managing /etc/fstab

Linux stores information about your local and remotely mounted filesystems in /etc/fstab. Open this file in the text editor of your choice. As you can see, different filesystems are configured on each line. A sample /etc/fstab might look like the following:

```
LABEL=/        /          ext3      defaults              1  1
LABEL=/boot    /boot      ext3      defaults              1  2
none           /dev/pts   devpts    gid=5,mode=620        0  0
none           /proc      proc      defaults              0  0
none           /dev/shm   proc      tmpfs                 0  0
/
```

Certification Objective 4.04-Access Control Lists and Other Security Attributes

Three of the methods for protecting individual files in Linux are based on file permissions, access control lists, and SELinux. File permissions are the standard method for security control, regulating access by user, group, and others. It's a basic prerequisite described in Chapter 1. With access control lists, file owners can regulate permissions for specific users. You can control SELinux from the command line; however, as described in Chapter 15, it's easier to control with the SELinux Management tool.

Exam Watch

Although ACLs and SELinux have just been added to the Red Hat Exam Prep guide, they have already been part of the course outline associated with both the RHCT and RHCE exams.

Access Control Lists

There was a time where users had read access to the files of all other users. But by default, users have permissions only in their own directories. Before you can configure ACLs, you'll need to set up execute permissions on the associated directories. For example, when I want to configure access to Donna's home directory, I first need to set appropriate permissions with the following command:

```
#
# chmod 701 /home/donna
```

Certification Summary

This chapter covers basic filesystem administration techniques on a Red Hat Enterprise Linux system. It also covers the different types of filesystems Linux uses, discusses how to mount them, and describes what mount options to use with them.

Creating a new filesystem means knowing how to create and manage partitions. Two excellent tools for this purpose are fdisk and **parted**.

You can automate this process for regular users with the automounter. Properly configured, it allows users to access shared network directories, removable media, and more through paths defined in /etc/auto.master.

Filesystems can be secured through attributes on individual files, access control lists, and settings available through Security Enhanced Linux.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 4.

Partitioning Hard Disks

? It's easiest if you can partition hard disks during the installation process.

? The fdisk tool can help you create and delete partitions, as well as change partition types.

? The **parted** tool can do everything that fdisk can do, and it can help you resize a partition.

Managing Filesystems

? Linux filesystems can be loosely defined as regular and journaling filesystems. While there are other filesystems available, this chapter describes the essential differences between the older ext2 and the current default ext3 filesystems.

? If you have the kernel source RPMs installed, you can review supported filesystems.

? A number of mount options are available for /etc/fstab. The **defaults** option sets up a partition as **rw** (read/write), suid (superuser ID and super group ID files allowed), dev (device files read), **exec** (binaries can be run), **auto** (automatic mounting), **nouser** (mountable only by root), and **async** (data is read asynchronously).

Filesystem Management and the Automounter

? Standard filesystems are mounted as defined in /etc/fstab.

? Portable filesystems such as CDs and USB keys are usually mounted automatically when installed.

? With the automounter, you can configure automatic mounts of removable media and shared network drives.

Access Control Lists and Other Security Attributes

? With Access Control Lists, you can allow specific users access to the files of your choice with the setfacl command.

?	With SELinux, Red Hat has implemented targeted control that protects network daemons, using fine-grained controls.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. While the topics in this chapter are "prerequisites," it is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

Partitioning Hard Disks

1. What fdisk command lists configured partitions from all attached hard drives?

   ____

2. What command from the **parted** prompt lists all created partitions?

   ____

3. If you've just installed a new hard drive and are configuring it in **parted**, what command do you need to run first?

   ____

4. After creating a swap partition, what command activates it?

   ____

Answers

1. The fdisk command that lists configured partitions from all attached hard drives is **fdisk -l**.
2. The command from the **parted** prompt that lists all created partitions is **print all**. The **print** command by itself just prints active partitions.
3. If you've just installed a new hard drive and are configuring it in **parted**, the command you need to run first is **mklabel**.
4. After creating a swap partition, the **swapon** ***partitionname*** activates it; just substitute the device file associated with the partition (such as /dev/sda10 or /dev/VolGroup00/LogVol03) for ***partitionname***.

Managing Filesystems

5. What is the primary advantage of a journaling filesystem such as ext3?

   ____

6. What can you run on /etc/samba/smb.conf to make sure even the administrative user can't delete it just using the **rm** command? The command must still allow the file to be readable. ?

____

7. What command formats /dev/sdb3 to the default Red Hat filesystem format? ?

____

Answers

[5.](#) The primary advantage of a journaling filesystem such as ext3 is faster data recovery if power is suddenly cut.

[6.](#) To make sure even the administrative user can't delete /etc/samba/smb.conf just using the **rm** command, run **chattr +i /etc/samba/smb.conf**. Other commands such as **chmod 000 /etc/ samba/smb.conf** provide some level of protection, but the file would no longer be readable.

[7.](#) The command that formats /dev/sdb3 to the default Red Hat filesystem format is **mkfs.ext3 /dev/sdb3**. Since ext3 is atop the list in /etc/filesystems, the **mkfs /dev/sdb3** command works as well.

Filesystem Management and the Automounter

8. To change the mount options for a local filesystem, what file would you edit? ?

____

9. If you've started the **autofs** daemon and want to read the list of shared NFS directories from the first.example.com computer, what automounter-related command would you use? ?

____

10. What daemon do you need to activate before using the automounter? ?

____

Answers

[8.](#) To change the mount options for a local filesystem, edit /etc/fstab.

[9.](#) If you've started the **autofs** daemon and want to read the list of shared NFS directories from the first.example.com computer, the automounter-related command you'd use to list those directories is **/etc/auto.net first.example.com**.

[10.](#) The daemon you need to activate before using the automounter is **autofs**.

Access Control Lists and Other Security Attributes

11. What setting do you need to add to /etc/fstab to make sure a filesystem is mounted with ACLs the next time you boot Linux? ?

____

12. What is the default policy associated with SELinux on Red Hat distributions?

Answers

11. The setting you need to add to /etc/fstab to make sure a filesystem is mounted with ACLs the next time you boot Linux is **acl**. For example, if you want to support ACL options, you can add the setting to the desired line in that file. The following would add ACL options to the directory labeled /home:

```
LABEL=/home /home ext3 defaults,acl 1 2
```

12. The default policy associated with SELinux on Red Hat distributions is *targeted*.

Lab Questions

The Red Hat exams are unique based on their reliance on labs and hands-on demonstrations. With these questions, you're practicing the skills you need on both Red Hat exams.

Lab 1

1. This lab assumes you have a new hard disk (or at least empty space on a current hard drive where you can add a new partition). You can simulate a new hard disk by adding appropriate settings to a VMware or Xen virtual machine. In this lab, you'll create a new partition using **parted**, format it, transfer the files currently on your /home (or if you don't have a lot of space, /tmp) directory to that partition, and revise /etc/fstab so the new partition is properly mounted the next time you boot Linux. ?

   If you have a limited amount of available space, dedicate only half of it to this lab and leave the other half empty for Lab 2.

Answers

1.

1. If you've been able to add a new hard drive, you should be able to review it from the **(parted)** prompt. But make sure to open the appropriate drive. For example, if it's the second SATA drive, do so with the **parted /dev/sdb** command.
2. Run the **print** command from the **(parted)** prompt. If it's a new drive, you'll see an "unrecognized disk label" message and can run **mklabel** to add an **msdos** label as described in the chapter. Otherwise don't run **mklabel**!
3. Make a note of available space in your partitions.
4. Create the new partition. The **mkpart** command provides prompts that help you define the new partition. If the partition is on a new hard drive, create a primary partition. Otherwise, you may be able to create only a logical partition.
5. Use the prompts to define the size of the partition from the start and ending MB location on the drive. As noted in the lab, make sure the size of the partition is half the available free space.
6. Run **print** again to confirm your changes. Make a note of the partition number. For example, if you've created partition 1 on /dev/sdb, the partition device file is /dev/sdb1.
7. Run **quit** to exit from **parted**. Run the partprobe command to make Linux reread the partition table (without rebooting).
8. Format the partition. Assuming you're using the default Red Hat format, use the **mkfs.ext3** ***partitionname*** command; substitute the device file for ***partitionname***.
9. Mount the new partition on a temporary directory; I often create a /test directory for this purpose. For the aforementioned partition, the command would be **mount /dev/sdb1 /test**.
10. Copy all of the files recursively from the directory that you're going to mount on the new partition. For example, if you're moving the files from the /home directory using the noted partitions, the command would be **cp -ar /home/* /test**.
11. Unmount **/test** from the new partition with a command like **umount /test**.
12. Mount the new partition such as /dev/sdb1 on the /home directory.
13. Review the results. Are the files you transferred on the new partition?
14.

Lab 2

2. In this lab, you'll add a new swap partition using the fdisk utility. Remember to make the partition work with the appropriate file type, and then format and activate it. Make sure it's properly included in /etc/fstab so this partition is used the next time you boot Linux.

Answers

2. In this lab, you'll add a new swap partition using the fdisk utility. Remember to make the partition work with the appropriate file type, format, and activate it. Make sure it's properly included in /etc/ fstab so this partition is used the next time you boot Linux.

    1. If you've completed Lab 1, you presumably have half the free space-from either an existing or a newly installed drive-still available.
    2. Use fdisk to open the drive with free space. You may need to be specific. The **fdisk -l** command can help you define the drive with free space, such as /dev/hdc. In that case, run the **fdisk /dev/hdc** command to edit that drive's partition table.
    3. Add a new partition using existing free space. From the fdisk prompt, the **p** command prints defined partitions, including the one you just created. Make sure to change the partition type; the **t** command from the fdisk prompt allows you to change the partition number you just created to the Linux swap system ID (82).
    4. Write your changes from fdisk; if you want to reread the partition table without rebooting, use the partprobe command.
    5. Format the new partition to the Linux swap filesystem: for example, if the new partition is on /dev/hdc3, you'd run the **mkswap /dev/hdc3** command.
    6. Once the format process is complete, you can immediately activate this partition with the **swapon /dev/hdc3** command.
    7. But that's not it. You need to make sure that swap partition is activated the next time you boot. To do so, you need to add information associated with that partition to /etc/fstab. One line that would work in this case is:

Lab 3

3. In this lab, you'll configure the automounter on your computer on an NFS connection, using two different methods. You'll need a second computer with Linux or Unix installed, and a shared NFS directory. You can use the shared NFS installation source created in Chapter 2 or any other shared NFS directory described in Chapter 10. A virtual machine such as a VMware computer qualifies as a second computer.

    1. Back up your current /etc/auto.master and /etc/auto.net configuration files.
    2. Open the /etc/auto.master configuration file in the text editor of your choice. Add or activate the command that applies the automounter to the /net directory.
    3. Open the /etc/auto.misc configuration file. Use the example shown in this file to create an NFS entry that points to the shared NFS directory on the second computer. For the purpose of this lab, I'll assume the name of the directory to test. Substitute accordingly.
    4. Restart the **autofs** server.
    5. Try your connection. Run the following command:

Answers

3. Configuring the automounter on a shared NFS directory is easier than it looks. Before you begin, make sure that you can mount the shared NFS directory from the remote computer. Resolve those problems first before beginning this lab. Refer to Chapter 2 on creating an NFS installation server or Chapter 10 on NFS for more information. If there's no problem with a source on an NFS server with an IP address of 192.168.30.4, you should be able to mount it locally. For example, you can mount a shared remote NFS /inst directory on an existing empty local /test directory as follows:

```
#
# mount -t nfs 192.168.30.4:/inst /test
```

Lab 4

4. In this lab, you'll configure access for the supervisor named Donna to the project.odt OpenOffice .org writer file in John's home directory, /home/john. Remember that you'll need to remount the appropriate partition, revise /etc/fstab, change permissions to /home/john, and set the ACL permissions to allow access by Donna.

Answers

4. As described in this chapter, you'll first have to change the mount of the filesystem with the /home directory to include ACL settings. For example, if /home isn't mounted separately from the top level root (/) directory and is part of /dev/VolGroup00/LogVol00, you can remount the filesystem with the /home directory with the following command:

```
#
# mount -o remount -o acl /dev/VolGroup00/LogVol00 /
```

Chapter 5: Package Management

After installation is complete, you still have administrative work to do. To customize the system to meet your needs, you may need to add or remove packages and more. To make sure you get the right updates, you need to know how to get your system working with the Red Hat Network (or the repository associated with your distribution). If you're satisfied with your configuration, you may want to use kickstart to automate future installations.

Certification Objective 5.01-The Red Hat Package Manager

One of the major duties of a system administrator is software management. New applications are installed. Services are updated. Kernels are patched. Without the right tools, it can be difficult to figure out what software is on a system, what is the latest update, and what applications depend on other software. Worse, you may install a new software package only to find it has overwritten a crucial file from a currently installed package.

Inside the Exam

Administrative Skills

The management of RPM packages is a fundamental skill for Red Hat administrators, so it's reasonable to expect to use the **rpm** and related commands on both parts of the RHCT or RHCE exam. While the requirements are listed in the Installation and Configuration part of each exam, these skills may also help you during the Troubleshooting and System Maintenance part of each exam. The Red Hat Exam Prep guide includes two specific references to the **rpm** command:

- Install and update packages using **rpm**.
- Properly update the kernel package.

And there's a related reference, as update tools such as yum and Pup are essentially front-ends to the **rpm** command:

- Configure the system to update/install packages from remote repositories using yum or Pup.

But whenever you install a new package, you need to know how to use the **rpm** command. You need to know how to use it to find the files you need. You need to know how the yum command helps manage dependencies, including any additional packages that might be required.

Finally, this chapter will also help you meet the Exam Prep guide requirement associated with Kickstart, which suggests that RHCEs must also be able to

- Configure hands-free installation using Kickstart.

The Red Hat Package Manager (RPM) was designed to alleviate these problems. With RPM, software is managed in discrete *packages*. An RPM package includes the software with instructions for adding, removing, and upgrading

those files. When properly used, the RPM system can back up key configuration files before proceeding with upgrades and removals. It can also help you identify the currently installed version of any RPM-based application.

RPMs and the **rpm** command are far from perfect, which is why it has been supplemented with the yum command. With a connection to a repository such as that available from the Red Hat Network or third-party "rebuilds" such as CentOS, you'll be able to use yum to satisfy dependencies.

What Is a Package?

In the generic sense, an RPM package is a container of files. It includes the group of files associated with a specific program or application, which normally includes binary installation scripts as well as configuration and documentation files. It also includes instructions on how and where these files should be installed and uninstalled.

An RPM package name usually includes the version, the release, and the architecture for which it was built. For example, the fictional penguin-3.4.5-26.i386 .rpm package is version 3.4.5, build 26, and the i386 indicates that it is suitable for computers built to the Intel 32-bit architecture.

On the Job

Many RPM packages are CPU-specific. You can identify the CPU type for your computer in the /proc/cpuinfo file. Some RPM packages with the noarch label can be installed on computers with all types of CPUs.

What Is an RPM?

At the heart of this system is the RPM database. Among other things, this database tracks the version and location of each file in each RPM. The RPM database also maintains an MD5 checksum of each file. With the checksum, you can use the **rpm -V *package*** command to determine whether any file from that RPM package has changed. The RPM database makes adding, removing, and upgrading packages easy, because RPM knows which files to handle and where to put them.

RPM also manages conflicts between packages. For example, assume you have two different packages that use configuration files with the same name. Call the original configuration file /etc/*someconfig.conf*. You've already installed package X. If you then try to install package Y, RPM packages are designed to back up the original /etc/*someconfig.conf* file (with a file name like /etc/*someconfig.conf.rpmsave*) before installing package Y.

On the Job

While RPM upgrades are supposed to preserve or save existing configuration files, there are no guarantees. It's best to back up all applicable configuration files before upgrading any associated RPM package.

Installing RPMs

There are three basic commands *that may* install an RPM. They won't work if there are dependencies (packages that need to be installed first). For example, if you haven't installed Samba and try to install the system-config-samba package, you'll get the following message (your version numbers may be different):

```
#
# rpm -i system-config-samba-*
e
error: Failed dependencies:
```

Certification Objective 5.02-More RPM Commands

The **rpm** command is rich with details, and learning to use this command can and does fill entire books, including the *Red Hat RPM Guide*, by Eric Foster-Johnson. All this book can do is cover some of the basic ways **rpm** can help you manage your systems. You've already read about how **rpm** can install and upgrade packages in various ways. Queries can help you identify what's installed, in detail. Validation tools can help you check the integrity of packages and individual files. You can use related tools to help you find the RPM that meets specific needs as well as a full list of what's already installed.

RPM Queries

The simplest RPM query verifies whether a specific package is installed. The following command verifies the installation of the Samba package:

```
#
```

Certification Objective 5.03-Managing Updates with Pup and the Red Hat Network

One key advantage of RHEL is access to the Red Hat Network (RHN). With the RHN, you can keep all of your registered systems up to date from one Web-based interface. You can even configure RHN to run commands remotely, on a schedule. Naturally, this can be a terrific convenience for remote administrators. But before any of this is possible, you'll need to register your system on the RHN.

If you're running a "rebuild" distribution such as CentOS-5, you don't need access to the RHN. If this applies to you, feel free to skim until reaching the "Updating with Pup" section.

As the RHN requires subscriptions, Red Hat does not have any public mirrors for RHEL updates. If you have a group of RHEL systems, it's possible to download the updates once and store them locally. You can then use those RPMs to update the other RHEL systems on your network. Red Hat facilitates this kind of communication with the RHN Proxy Server and Satellite Server products.

Exam Watch

While the Red Hat Network is covered per the public syllabus for the prep courses for both the RHCT and RHCE, it is not included in the Red Hat Exam Prep guide. And I don't think you'll have Internet access during the Red Hat exams.

While you can configure your own local proxy of updates, it won't come with Red Hat support. I've written a guide to this process, *Linux Patch Management*, published by Prentice-Hall.

RHN Registration

Before you can administer your system on the RHN, you have to register. You'll need either a registration code associated with your subscription or available entitlements for your RHN account.

You may have already registered during the First Boot process described in Chapter 2. But if you've installed the system in text-mode, or using kickstart (common on many servers), you may not have gone through the First Boot process, at least the version that supports registration. Naturally, you can still register after installation. This section covers text-mode registration, which is how most administrators work with servers. If you want to register your system from the GNOME Desktop Environment, you're prompted to do so during the first update. Choose Applications | System Tools | Package Updater, and follow the wizard. If you don't see the wizard, your system is already registered (or you're using a rebuild distribution).

To register from the command line, take the following steps:

1. Run **rhn_register** from the command line.
2. You'll see a Setting Up Software Updates window. If you need more information about the RHN, select Why Should I Connect To RHN; otherwise, select Next to continue.
3. You'll see a screen where you can enter your login information for the RHN. If you don't have an RHN account, select Create A New Login, and follow the instructions there before moving on to step 4. Otherwise,

enter your RHN account information and select Next.

4. Now you can choose whether to register a system profile. First, you can choose whether to send basic hardware information about your system; make a decision and select Next to continue.

5. Next, you can choose to include a list of installed packages, which helps the RHN check whether you need software and security updates. Make any desired changes and select Next to continue.

6. Finally, you can choose whether to send your system profile to the RHN. If you click Cancel, the tool stops, and your system is not registered. I assume you want to register; if so, select Next to continue.

7. Your system attempts to contact the RHN server (or possibly your RHN Satellite Server). After additional prompts, you should see a message that you've successfully registered your system with the RHN.

Updating with Pup

Red Hat has adapted the Package Updater, also known as Pup, to manage updates for RHEL systems. If you're properly registered on the RHN, you can use Pup to list available updates and download them as needed. It runs only in the GUI. To start it, choose Applications | System Tools | Software Updater. As shown in Figure 5-1, it's a simple tool; it lists only those packages for which later versions are available, which you can deselect as desired before newer packages are downloaded and installed. If you select Update Details, you can review the version numbers of the packages being changed.

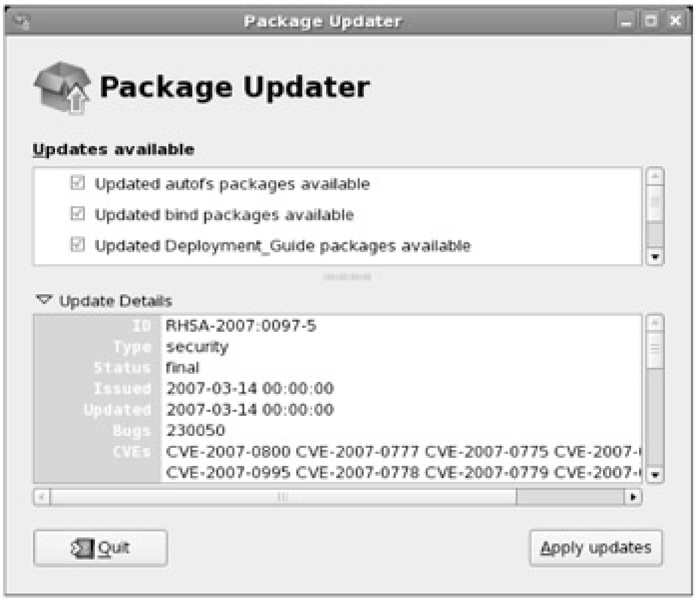

Figure 5-1: Pup, the Package Updater

On the Job

If you haven't already registered your system, starting Pup first starts the Red Hat Network Registration steps just described.

When you've made your selections, click Apply Updates.

Automatic Dependency Resolution

Red Hat has incorporated dependency resolution into the update process. Through RHEL 4, this was done with **up2date**. Red Hat has now incorporated yum into RHEL 5. The yum command uses subscribed RHN channels and

any other repositories that you may have configured in the /etc/yum.repos.d directory.

Before yum and **up2date**, dependencies were a serious annoyance. For example, if you hadn't installed the Samba RPM and tried to install the system-config-samba RPM, the installation would fail with a message like this:

```
e
error: Failed Dependencies:
```

Certification Objective 5.04-Adding and Removing RPM Packages with yum and pirut

The yum command makes it easy to add and remove software packages to your system. It maintains a database regarding the proper way to add, upgrade, and remove packages. This makes it relatively simple to add and remove software with a single command.

The Basics of yum

If you want to know all that you can do with yum, run the command by itself:

```
#
# yum
 L
 Loading "rhnplugin" plugin
 L
 Loading "installonlyn" plugin
Y
You need to give some command



u
usage: yum [options] <grouplist, localinstall, groupinfo,
 l
 localupdate, resolvedep, erase, deplist, groupremove,
m
makecache, upgrade, provides, shell, install, whatprovides,
```

Certification Objective 5.05-Using Kickstart to Automate Installation

Kickstart is Red Hat's solution for an automated installation of Red Hat. All the questions asked during setup can be automatically supplied with one text file. You can easily set up nearly identical systems very quickly. Kickstart files are useful for quick deployment and distribution of Linux systems.

Inside the Exam

Kickstart and the Red Hat Exams

For the Troubleshooting and System Maintenance portion of the exam, it's possible that the exam proctor might configure your computer using a customized kickstart configuration file. (It's possible that she might use another method such as a VMware snapshot.) The file might be local or it might be stored on the server. Understanding kickstart is a very useful skill that can help you install Linux on a number of different computers simultaneously. You can start the process and walk away. The options are rich and varied. The Red Hat Exam Prep guide suggests that you know how to configure "hands-free installation using Kickstart."

Whether or not you see kickstart labs on your exam, understanding how it works is an important administrative skill that I think Red Hat should include on its exams. You can check the bug status for yourself at https://bugzilla.redhat.com.

Kickstart Concepts

There are two methods for creating the required kickstart configuration file:

- Start with the anaconda-ks.cfg file from the root user's home directory, /root.
- Use the graphical Kickstart Configurator, accessible via the **system-config-kickstart** command.

The first option lets you use the kickstart template file created for your computer by Anaconda, the Red Hat Enterprise Linux installation program. The second option, the Kickstart Configurator, is discussed in detail later in this chapter.

If you're installing RHEL on a number of computers, the anaconda-ks.cfg file is handy. You can install RHEL the way you want on one computer. You can then use the anaconda-ks.cfg file from that computer as a template to install RHEL on the other identical computers on your network. If the computers aren't identical, you can customize each anaconda-ks.cfg file as required for elements such as a different hard disk size, host name, and so on.

Setting Up a Kickstart USB

Once the kickstart file is configured, the easiest way to use it is through the RHEL installation USB drive. Similar steps can be used if your system has a floppy drive. To do so, follow these basic steps:

1. Configure and edit the anaconda-ks.cfg file as desired. I'll describe this process in more detail shortly.
2.

Insert and mount the image file for an installation boot CD or USB drive. These options are described in more detail in Chapter 2. If the drive doesn't mount automatically, you can then mount the drive with a command such as **mount /dev/cdrom /mnt**.

3.

If you mounted a CD, insert a USB drive. Find the associated device with the **fdisk -l** command. Make a note of the device, such as /dev/sdb. You'll need it shortly.

4.

If it doesn't mount automatically (as it should if you're in the GNOME desktop), mount it with a command such as **mount /dev/sdb /net**.

5.

Copy the kickstart file to the top level directory of the USB drive. Make sure the file name on the USB drive is ks.cfg.

6.

You should now be ready to try out the installation boot USB on a different computer. You'll get to try this again shortly in an exercise.

7.

When you try the boot USB, you can start the kickstart installation from the aforementioned USB drive, possibly with the following command (some trial and error may be required):

Certification Summary

This chapter covered the management of RPM packages. You learned how to add, remove, and upgrade packages, and how to add updates-locally and remotely. It's important to upgrade kernels by installing them, side by side with currently working kernels.

You also learned how to query packages, examine to which package a file belongs, the steps necessary to validate a package signature, and how to find the current list of installed RPMs. You also read about installing and building source RPMs.

If you have an RHEL subscription, you can keep your system up to date through the RHN. You can install more packages from the RHN or configured repositories using yum. Alternatively, you can connect to the same repositories using tools such as pirut. In either case, yum provides automatic dependency resolution, which simplifies the installation and update process.

You can automate your entire installation with kickstart. Every RHEL system has a kickstart template file in the /root directory, which you can modify and use to install RHEL on other systems automatically. Alternatively, you can use the GUI Kickstart Configurator to create an appropriate kickstart file.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 5.

The Red Hat Package Manager

- ? The RPM database tracks where each file in a package is located, its version, and much more.
- ? Verifying an installed package confirms the integrity based on the RPM database.
- ? The Install mode of RPM installs RPM packages on your system; a newly installed kernel is loaded side by side with a previously installed kernel.
- ? The Upgrade mode of RPM replaces the old version of the package with the new one.

More RPM Commands

- ? The **rpm -e** command (erase) removes a package from your system.
- ? The **rpm** command query mode (**-q**) determines whether packages are installed on your system or files are associated with a particular package.
- ? Source RPMs, as the name indicates, contain the source code used to build architecture-specific packages.
- ? The spec file loaded in /usr/src/redhat/SPECS/packagename.spec controls the way a package is built and what actions are performed when it is installed or removed from a system.
- ? Run **rpmbuild -ba** ***packagename*****.spec** to build your binary and source RPM.

Managing Updates with Pup and the Red Hat Network

- ? Before connecting to the RHN, you need to register your system.
- ? The Package Updater, Pup, can help you keep systems up to date.
- ? With automatic dependency resolution, yum and the RHN help install dependencies along with desired packages.

? The RHN can help you manage subscribed systems remotely using a Web-based interface.

Adding and Removing RPM Packages with yum and pirut

? The yum command can help install a group of packages from the RHN or repositories configured in the /etc/yum.repos.d directory.

? The pirut tool can help you add and remove packages and complete package groups from your system.

Using Kickstart to Automate Installation

? Kickstart is Red Hat's solution for an automated simultaneous installation on several computers.

? Kickstart installations can be configured to take installation files from a CD-ROM, a local drive, an NFS, an FTP, or an HTTP server.

? There are two ways to create a kickstart file: from the configuration when you installed Linux as documented in the /root/anaconda-ks.cfg file, or from the GUI Kickstart Configurator.

PREV NEXT

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

The Red Hat Package Manager

1. What command would you use to install the penguin-3.26.i386.rpm package, with extra messages in case of errors? The package is on the local directory. ?

   ____

2. What command would you use to upgrade the penguin RPM with the penguin-3.27.i386.rpm package? The package is on the ftp.remotemj02.abc server. ?

   ____

3. If you've downloaded a later version of the Linux kernel to the local directory, and it's kernel-2.6.19.el5.i386.rpm, what's the best way to make it a part of your system? ?

   ____

Answers

1. The command that installs the penguin-3.26.i386.rpm package, with extra messages in case of errors, from the local directory, is

```
#
# rpm -iv penguin-3.26.i386.rpm
```

2. The command that upgrades the aforementioned penguin RPM with the penguin-3.27.i386.rpm package from the ftp.remotemj02.abc server is

```
#
```

3. If you've downloaded a later version of the Linux kernel to the local directory, and it's kernel-2.6.18-8.4.4.el5.i386.rpm, the best way to make it a part of your system is to install it-and not upgrade the current kernel. Kernel upgrades overwrite existing kernels. Kernel installations allow kernels to exist side by side; if the new kernel doesn't work, you can still boot into the working kernel. So you'd use a command like this:

```
#
# rpm -iv kernel-2.6.18-8.4.4.el5.i386.rpm
```

More RPM Commands

4. What command lists all installed RPMs? What about the RPMs installed when you first installed the local system?

____

____

5. Assume you have the rpm-build RPM installed. When you install a source RPM, in what directory (and subdirectories) will you find key files?

____

6. If you've downloaded an RPM from a third party and called it third.i386.rpm, how would you validate the associated package signature?

____

Answers

4. The command that lists all installed RPMs is

```
#
# rpm -qa
```

5. When you install a source RPM, the directory in which key files are stored is /usr/src/redhat. One important file is the spec file associated with the RPM in the /usr/src/redhat/SPECS directory.

6. If you've downloaded an RPM from a third party, call it third.i386.rpm, you'll first need to download and install the RPM-GPG-KEY file associated with that repository. You can then validate the associated package signature with a command like (note the uppercase **-V**):

```
#
# rpm -V third.i386.rpm
```

Managing with Pup and the Red Hat Network

7. If you want official updates for RHEL, to where should you connect your RHEL system?

____

Answers

7. If you want official updates for RHEL, you should connect your RHEL system to the Red Hat Network.

Adding and Removing RPM Packages with yum and pirut

8. Name at least two tools that can help you download and install updates from the RHN.

____

____

9. What [yum](#) command installs the latest Linux kernel?

?

Answers

[8.](#) Three tools can help you download and install updates from the RHN: Pup, [pirut](#), and [yum](#).

[9.](#) This [yum](#) command installs a later available version of the Linux kernel:

```
#
# yum install kernel
```

Using Kickstart to Automate Installation

10. You're using the Kickstart Configurator to create a ks.cfg file for several computers. Interpret the following directive:

?

```
p
part /var --size 1000 --grow yes
```

11. If your kickstart installation file is on the local CD, and you boot from the USB drive, what would you type at the **boot:** prompt to start the kickstart installation?

?

---

12. If your kickstart installation file is on the local hard drive in /dev/sda7, on the top-level directory on that partition, and you boot from the USB drive, what would you type at the **boot:** prompt to start the kickstart installation?

?

Answers

[10.](#) The following directive in ks.cfg configures a partition for the /var directory, of at least 1000MB, but growable-which means it can take up the remaining free space.

```
p
part /var --size 1000 --grow yes
```

[11.](#) If your kickstart installation file is on the local CD, and you boot from the USB drive, type the following at the **boot:** prompt to start the kickstart installation:

```
l
linux ks=cdrom:/ks.cfg
```

[12.](#) If your kickstart installation file is on the local hard drive in /dev/sda7, on the top-level directory, and you boot from the USB drive, type the following command at the **boot:** prompt to start the kickstart installation:

```
l
linux ks=hd:sda7:/ks.cfg
```

Lab Questions

The Red Hat exams are unique based on their reliance on labs and hands-on demonstrations. With these questions, you're practicing the skills you need on both Red Hat exams.

Lab 1

1. In this lab, you'll examine what happens when you update a kernel RPM by installing it side by side with an existing kernel. If a newer kernel is not available, the kernel-xen package, or even an older kernel, will serve the purpose for this lab. Just remember, if you don't want the kernel that you install during this lab, make sure to remove the package properly from your system.
   1. Make a copy of your existing GRUB configuration file, /boot/grub/grub.conf. Print it out or copy it to your home directory.
   2. Make a copy of the current file list in the /boot directory. One method uses the **ls /boot > bootlist** command, which writes the file list to the bootlist file.
   3. If a newer kernel is available, and you're connected to the RHN or another appropriate repository, run the following command:

Answers

1. This lab is somewhat self-explanatory and is intended to help you explore what happens when you properly install a new kernel RPM. As with other Linux distributions, when you install (and do not use upgrade mode) for a new kernel, two areas are affected.

   The new kernel is added as a new option in the GRUB configuration menu. Unless you've installed an older kernel, the default boot option does not change. However, when you reboot, you'll be able to select the new kernel from the GRUB menu.

   When you review the /boot directory, all of the previously installed boot files should be there. The new kernel RPM should add matching versions of all of the same files-with different revision numbers (unless it's a Xen-based kernel).

   To keep this all straight, it helps if you made copies of the original versions of the GRUB configuration file and the file list in the /boot directory.

Lab 2

2. Generally, the only reason you need to install a kernel source RPM is if you absolutely need to recompile the kernel. Drivers can often be compiled using the kernel-devel package. However, if you need the source code, you may have to download it directly from the repository associated with your distribution. If you're running RHEL 5, that's available through your RHN subscription or ftp .redhat.com. If you're running Fedora Core, you'll need to activate the applicable source repository. If you're running a rebuild, you may need to download the kernel-*versionnum*.src.rpm directly from the repositories associated with that rebuild.

   While "rebuilds" are supposed to use the same source code as Red Hat, there is no guarantee as such. I've run into trouble when mixing the source code released with different rebuild distributions. So it's best if you download the source code from the associated repositories.

   Once you download the source RPM, you can install it with the **rpm -ivh kernel-*versionnum* .src.rpm** command, but that just starts the process of unpacking the source code. You'll need the rpmbuild command, available from the rpm-build RPM. You can then navigate to the /usr/src/redhat/SPECS directory, and use the **rpmbuild -bb** command to unpack the source code to different directories in the /usr/src/redhat/BUILD/ directory tree.

   But wait, the kernel source code already seems to be in the /usr/src/kernels directory, in a subdirectory named for the kernel version. However, this source code is not complete; it's intended only for building drivers, and if you want to recompile a kernel, you still need to apply the **rpmbuild -bb** command to the kernel-2.6.spec file in the /usr/src/redhat/SPECS directory.

Answers

2. This lab can help you prepare for Chapter 8, where you'll recompile the Linux kernel. Red Hat no longer provides a binary RPM for the source code. The process for installing a kernel source code RPM is subtly different from other source code RPMs, as it loads the source code into unique directories.

Lab 3

3. This lab may not be possible unless updates are available from your repository or the RHN. In this lab, you'll examine what happens when you run an update to upgrade to newer versions of packages available for new features, to address security issues, and more. Before you start, run the following command to clear the cache, so you get the full set of messages:

```
#
# yum clean all
```

Answers

3. This lab is intended to help you examine what the **yum update** command can do. It's the essential front end to other update tools, namely Pup. As you can see from the update.txt file created in this lab, the messages show you how yum looks for all newer packages from configured repositories or the RHN, downloads their headers, and uses them to check for dependencies that also need to be downloaded and installed.

Lab 4

4. In this lab, you'll get a chance to use the configuration for your current system to kickstart an installation of a second system. Ideally, you'll have a VMware or Xen virtual machine available for the process, with an identical amount of free space and hardware as the current system. Otherwise, this lab may not work.

   Open the anaconda-ks.cfg file in the current installation of RHEL. Remove the comments as appropriate from the directives associated with partitions and filesystems. Configure an installation boot CD or USB key, depending on what you can boot from your system. Copy the revised kickstart file to ks.cfg, and write it to appropriate media, even a floppy drive if available.

   Make sure the same source you used for the original installation (network, hard drive, CD/DVD) is still available. Boot the new system to test the installation.

Answers

4. For a lab like this, it's critical that you have a *second* system in which you don't mind losing all data. VMware and Xen are excellent options for this purpose. If successful, the kickstart installation you create and run will erase all data on that second system (unless specially configured). If space is limited, you can certainly delete the virtual machine files associated with this second system after installation.

PREV NEXT

Chapter 6: User Administration

Overview

For the Red Hat exams, the skills you learn in this chapter are important for the Installation and Configuration portion of each exam. As described in the Red Hat Exam Prep guide, you need to know how to manage accounts and set up the user environment.

As part of learning how to set up the user environment, you will learn how to set up the Linux startup shell configuration scripts so that users' sessions are configured according to your (and their) requirements. You will learn how to create and implement policies for managing disk usage-by user or by group. Special groups can help users share files securely.

There are different ways to secure your system and network. The PAM (Pluggable Authentication Modules) system lets you configure how users are allowed to log in or access different services. The Network Information Service (NIS) and the Lightweight Directory Access Protocol (LDAP) can provide a common database of authentication and configuration files for your network.

Inside the Exam

This chapter addresses four items as listed in the Red Hat Exam Prep guide. Three are associated with the Installation and Configuration section of the RHCT exam requirements:

- Attach system to a network directory service, such as NIS or LDAP.
- Add and manage users, groups, and quotas, and File Access Control Lists.
- Configure filesystem permissions for collaboration.

(File Access Control Lists were already addressed in Chapter 4.)

Remember that if you're studying for the RHCE, you have to know all the RHCT requirements. In addition, one item in this chapter is associated with the Installation and Configuration section of the RHCE exam requirements:

- Use PAM to implement user-level restrictions.

When you take the Red Hat exams, as long as you don't cheat, it generally does not matter how you come to a solution. For example, you get the same credit whether you add users by directly editing /etc/passwd using commands such as **useradd** or by using GUI tools such as the Red Hat User Manager (**system-config-users**). As in the real world, it is the results that matter.

Certification Objective 6.01-User Account Management

You need to know how to create and configure users for the Red Hat exams. This means that you need to know how to configure the environment associated with each user account-in configuration files and in user settings. You also need to know how to specify the configuration files associated with the default bash shell. Finally, you need to know how to limit the resources allocated to each user through quotas. These requirements are all explicitly cited in the Red Hat course outlines associated with the RHCT exam and are applicable to both exams.

If you've installed RHEL 5 via kickstart or in text mode, the default Red Hat installation gives you just a single login account: root. You should set up some regular user accounts. You may have already done so through the First Boot process described in Chapter 2. Even if you're going to be the only user on the system, it's a good idea to create at least one nonadministrative account to do your day-to-day work. Then you can use the root account only when it's necessary to administer the system. Accounts can be added to Red Hat Enterprise Linux systems using various utilities, including application of the vi text editor (and related specialized commands) on password configuration files (the manual method), the **useradd** command (the command line method), and the Red Hat User Manager utility (the graphical method).

Exam Watch

As discussed earlier, it's faster to log in as root (and not just the superuser). While you'll be doing most of the work on the Red Hat exams as root, it's quite possible that you'll be asked to create accounts for regular users (and groups) to configure a workstation.

User Account Categories

There are three basic types of Linux user accounts: administrative (root), regular, and service. The administrative root account is automatically created when you install Linux, and it has administrative privileges for all services on your Linux computer. A cracker who has a chance to take control of this account can take full control of your system.

Nevertheless, it is sometimes appropriate to log in as an administrator (that is, as the root user), such as during most of the Red Hat exams. Red Hat Enterprise Linux builds in safeguards for root users. Log in as the root user, and then run the **alias** command. You'll see entries such as this,

```
a
alias rm='rm -i'
```

Certification Objective 6.02-The Basic User Environment

Each user on any Red Hat Enterprise Linux system has an environment when logged on to the system. The environment defines directories where Linux looks for programs to run, the look of the login prompt, the terminal type, and more. This section explains how you can configure the default environment for your users.

Home Directories and /etc/skel

By default, when you create a new user, a default set of configuration files is created in the user's home directory. These defaults are hidden files stored in the /etc/skel directory.

Home Directory

The home directory is where a user starts when he first logs on to a RHEL system. The home directory for most users is /home/*username,* where *username* is the user's login name. Every user should normally have write permission in his own home directory, so each user is free to read and write his own files. Later in this chapter, you'll learn how to configure disk quotas, so users don't take more than their fair share of disk space.

/etc/skel

The /etc/skel directory contains default environment files for new accounts. The **useradd** command and the Red Hat User Manager copy these files to the home directory when you create a new account. The contents of /etc/skel may vary depending on what you have installed. Standard files from my copy of /etc/skel are described in Table 6-5.

Table 6-5: Default Home Directory Files from /etc/skel

File	Purpose
.bashrc	This basic bash configuration file may include a reference to the general /etc/bashrc configuration file. Can include commands you want to run when you start the bash shell, appropriate for aliases such as **rm='rm -i'**.
.bash_logout	This file is executed when you exit a bash shell and can include commands appropriate for this purpose, such as commands for clearing your screen.
.bash_profile	Configures the bash startup environment. Appropriate place to add environment variables or modify the directories in your PATH.
.gtkrc	Adds the Bluecurve theme for the default Red Hat GUI.
.kde	This directory includes autostart settings for the K Desktop Environment. Not added to /etc/ skel and not copied to user home directories if you haven't installed KDE on this computer.

If you've installed more than the default software on RHEL, you may see additional configuration files in the /etc/skel directory. For example, if you've installed the appropriate RPMs, you may see configuration files associated with emacs and the z shell (zsh) in this directory.

On the Job

Linux includes many hidden files that start with a dot (.). To list these files, run the *ls -a* command. For example, if you want to list all of the files in the /etc/skel directory, run the *ls -a /etc/skel* command.

As the system administrator, you can edit these files or place your own customized files in /etc/skel. When new users are created, these files are propagated to the new users' home directories.

On the Job

Adding files to /etc/skel may be a convenient way to distribute files such as corporate policies to new users.

Window Manager Configuration Files

RHEL comes with several window managers. At some point you will want to configure one or more of them for use on your system. In the past, window manager configuration files were stored in the /etc/X11/*windowmanager* directory, where *windowmanager* is the name of the specific window manager. This directory still includes twm (Tom's Window Manager). The X Display Manager is no longer available for RHEL 5. The GNOME and KDE Display Managers are stored in their own directories, /etc/gdm and /usr/share/ config/kdm.

PREV NEXT

Certification Objective 6.03-Shell Configuration Files

All system-wide shell configuration files are kept in the /etc directory. These files are bashrc, profile, and the scripts in the /etc/profile.d directory. These files and scripts are supplemented by hidden files in each user's home directory, as just described. Let's take a look at these files.

/etc/bashrc

The /etc/bashrc file is used for aliases and functions, on a system-wide basis. Open this file in the text editor of your choice. Read each line in this file. Even if you don't understand the programming commands, you can see that this file sets the following bash shell parameters for each user. For example:

- It assigns a value of umask, which creates the default permissions for newly created files. It supports one set of permissions for root and system users (with user IDs below 100), and another for regular users. (Officially, RHEL reserves all user IDs above 500 for regular users but that is not reflected in /etc/bashrc.)
- It assigns a prompt, which is what you see just before the cursor at the command prompt.
- It includes settings from *.sh files in the /etc/profile.d/ directory.

The settings here are called by the .bashrc file in each user's home directory. The settings are supplemented by the .bash_history and .bash_logout files in each user's home directory.

/etc/profile

The /etc/profile file is used for system-wide environments and startup files. The following is the profile script from my copy of the operating system. The first part of the file sets the PATH for searching for commands, adding more directories using the **pathmunge** command. (Unless you use the Korn shell, ignore the **ksh workaround** stanza.) Then it sets the PATH, **USER**, **LOGNAME**, **MAIL**, **HOSTNAME**, **HISTSIZE**, and **INPUTRC** variables and finally runs the scripts in the /etc/profile .d directory. You can check the current value of any of these variables with the **echo $*variable*** command.

```
#
# /etc/profile


#
# System wide environment and startup programs, for login setup
#
# Functions and aliases go in /etc/bashrc


 p
 pathmunge () {

  if ! echo $PATH | /bin/egrep -q "(^|:)$1($|:)" ; then
```

```
    if [ "$2" = "after" ] ; then
        PATH=$PATH:$1
    else
        PATH=$1:$PATH
    fi
  fi
}
}

#
# ksh workaround
i
```

Certification Objective 6.04-Setting Up and Managing Disk Quotas

Quotas are used to limit a user's or a group of users' ability to consume disk space. This prevents a small group of users from monopolizing disk capacity and potentially interfering with other users or the entire system. Disk quotas are commonly used by Internet Service Providers (ISPs), by Web hosting companies, on FTP sites, and on corporate file servers to ensure continued availability of their systems.

Without quotas, one or more users can upload files on an FTP server and occupy all free space on a partition. Once the affected partition is full, other users are effectively denied upload access to the disk. This is also a reason to mount different filesystem directories on different partitions. For example, if you only had partitions for your root (/) directory and swap space, someone uploading to your computer could fill up all of the space in your root directory (/). Without at least a little free space in the root directory (/), your system could become unstable or even crash.

You have two ways to set quotas for users. You can limit users by inodes or by kilobyte-sized disk blocks. Every Linux file requires an inode. Therefore, you can limit users by the number of files or by absolute space. You can set up different quotas for different filesystems. For example, you can set different quotas for users on the /home and /tmp directories if they are mounted on their own partitions.

Limits on disk blocks restrict the amount of disk space available to a user on your system. Older versions of Red Hat Linux included LinuxConf, which included a graphical tool to configure quotas. As of this writing, Red Hat no longer has a graphical quota configuration tool. Today, you can configure quotas on RHEL only through the command line interface.

On the Job

Learn to focus on command line tools. Red Hat used to make LinuxConf available as a graphical and console tool for a number of system administration functions, including quotas. While Red Hat may eventually create another GUI quota manager, don't count on it.

Quota Settings in the Kernel

By default, the Linux kernel as configured by Red Hat supports quotas. However, if you install and compile a new kernel from a remote source, you should make sure that this feature is active. The basic kernel configuration is stored in the /boot directory. For the default RHEL system, you'll find the configuration in the config-*versionnumber* file. If you've configured a custom kernel file, you'll find it listed under a different name.

To verify that quotas are enabled in the default kernel, run the following command (the shell substitutes the actual version number of the kernel for **`uname -r`**):

```
#
# grep CONFIG_QUOTA /boot/config-`uname -r`
```

Certification Objective 6.05-Creating and Maintaining Special Groups

One major difference between Red Hat Enterprise Linux and non-Red Hat Linux or Unix distributions is how new users are assigned to groups. A Linux group allows its members to share files. Unfortunately, that also means everyone in the same primary group has access to the home directories of all other group members. Users may not always want to share the files in their home directories with others. For example, if you're setting up an ISP, your users pay for their privacy.

On the other hand, RHEL gives each user a unique user ID and group ID in /etc/ passwd. This is known as the *user private group* scheme. Users get exclusive access to their own groups and don't have to worry about other users reading the files in their home directories.

On the Job

There are other ways to provide access to other users, as discussed in Chapter 4's "Access Control Lists."

Standard and Red Hat Groups

Traditionally, users are assigned to one or more groups such as users in /etc/group. For example, you might configure accgrp for the accounting department and infosys for the information systems department in your company.

If you have access to one of these other versions of Unix or Linux, check the third and fourth fields in /etc/passwd. Many users will have the same fourth field, which represents their *primary* group. Then, when you create a new user, each account receives a unique user ID but shares the same group ID with other users in the acct group. Users can also belong to other groups.

In RHEL, each user gets her own special private group by default. As you probably noticed earlier, user IDs and group IDs by default start at 500, match, and proceed in ascending order.

By default in RHEL, all regular users have a umask of 0002. If you are coming from a traditional Unix environment, you may be concerned. With the traditional user/group scheme, any member of that user's primary group will automatically have write access to any file that the user creates in his home directory.

This is the advantage behind the user private group scheme. Since every user account is the only member in its own private group, having the umask set to 0002 does not affect file security. This provides advantages for systems such as ISPs, where you don't want users to have access to each other's files.

Shared Directories

Most people work in groups. They may share files. You can give a group of users access to a specific user's home directory or you can set up a shared directory for a group.

When you configure a shared directory, you can set up a group owner and then add the users to that group through the /etc/group configuration file. When you set the group ID bit (SGID) on this directory, any file created in this directory inherits the group ID. Assuming you have set appropriate permissions, all members of this group can then access files in that directory.

There are several basic steps required to create a useful shared directory. For example, assume you want to set up a

shared directory, /home/accshared, for the accountants in your organization. To set this up, take the following steps:

1. Create the shared directory:

Certification Objective 6.06-Pluggable Authentication Modules

RHEL uses the Pluggable Authentication Modules (PAM) system to check for authorized users. PAM includes a group of dynamically loadable library modules that govern how individual applications verify their users. You can modify PAM configuration files to suit your needs.

Exam Watch

PAM modules are documented in the /usr/share/doc/pam-versionnumber/txts directory. For example, the functionality of the pam_securetty .so module is described in the README .pam_securetty file.

PAM was developed to standardize the user authentication process. For example, the login program uses PAM to require usernames and passwords at login. Open the /etc/pam.d/login file. Take a look at the first line:

```
a
auth [user_unknown=ignore success=ok ignore=ignore default=bad] \
```

Certification Objective 6.07-Network Authentication Configuration: NIS and LDAP

By default, access to a Linux computer requires a valid username and password. One problem with a large network of Linux systems is that "normally," each user requires an account on every Linux computer.

The two services that allow you to set up one centrally managed database of usernames and passwords for Linux and Unix computers are NIS and LDAP. With each of these services, you can maintain one password database on an NIS or LDAP server and configure the other systems on the network as clients. When a user logs into an NIS or LDAP client, that system first checks its local password file, usually /etc/passwd. If it can't find your username, it looks up the corresponding file on the server.

Exam Watch

In the Red Hat Exam Prep guide, the only requirement is to be able to connect a client to a network directory service, such as NIS or LDAP. As of this writing, the prep course outline for the RHCE (RH300) no longer includes NIS server configuration requirements. I therefore focus on NIS and LDAP clients in this section.

First, I'll show you how you can configure NIS and LDAP clients using the command line interface and then use the Red Hat Authentication Configuration tool.

NIS Client Configuration

It's fairly simple to configure an NIS client on a network. Assuming you have an NIS server, you need to do three things. First, specify the server and domain name in /etc/yp.conf. Next, make sure the ypbind client service starts the next time you boot Linux. Finally, make sure the /etc/nsswitch.conf file looks to the NIS service for at least the username and password database.

The change to the /etc/yp.conf configuration file is simple. All you need is a command such as the following, which specifies the name of the NIS domain as **nisdomain**, and the name of the NIS server as **enterprise5a**:

```
d
domain nisdomain server enterprise5a
```

Certification Summary

You can have great control over how your Linux installation is set up and configured. You can configure users and groups and control almost all aspects of the user environment. Any variables or system-wide functions you may need to run can be kept in the /etc/bashrc or /etc/profile script.

You can set up quotas to limit the user's disk usage. You can set up one quota per partition and set soft and hard limits for users. With grace periods, you can set up a soft limit to give users an appropriate warning.

By default, Red Hat Enterprise Linux assigns unique user and group ID numbers to each new user. This is known as the user private group scheme. This scheme allows you to configure special groups for a specific set of users. The users in the group can be configured with read and write privileges in a dedicated directory, courtesy of the SGID bit.

RHEL 5 includes powerful tools for securing critical commands, using Pluggable Authentication Modules (PAM). You can use centralized account management with an NIS or LDAP service.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 6.

User Account Management

? After installation, your system may have only a single login account: root. For most installations, you'll want to create more accounts, if only for everyday use.

? Accounts can be added by directly editing /etc/passwd or with the **useradd** command. The advantage of **useradd** is that it automatically adds the new home directory as well as configuration files from /etc/skel.

? Accounts can be added with the Red Hat User Manager tool. You can also use this tool or related commands such as chage and usermod to configure an account further with parameters such as a password lifetime or a time limit on the account.

? Discourage the use of shared accounts, where several people use a single account. Shared accounts are almost always unnecessary, and they are easily compromised.

? If you're using the Network File System (NFS), it can help establish user accounts with the same UID across systems.

? The Network Information System (NIS) can serve the same purpose by establishing one database for all systems on your network.

The Basic User Environment

? Each user on your system has an environment when logged on to the system.

? The home directory for each login account is the initial directory in which users are placed when they first log on. They start with hidden files configured in the /etc/skel/ directory.

Shell Configuration Files

? All system-wide shell configuration files are kept in the /etc directory.

? /etc/profile is the system-wide startup shell script for bash users.

? All users have hidden shell configuration files in their home directories.

Setting Up and Managing Disk Quotas

? Quotas are used to limit a user's or a group of users' ability to consume disk space.

? Quotas are set on specific filesystems mounted to standard Linux formats.

? Quota support must be enabled in the kernel. By default, quotas are enabled in RHEL kernels.

? Quotas have soft limits and hard limits. If both soft and hard limits are set, a user can exceed his or her soft limit for a modest period of time.

? Users and groups may never exceed their hard quota limits.

Creating and Maintaining Special Groups

? Red Hat's user private group scheme configures users with their own unique user and group ID numbers.

? With appropriate SGID permissions, you can configure a shared directory for a specific group of users.

? Setting the SGID bit ensures that all files created in a shared directory belong to the correct group.

? Setting the SGID bit is easy; use chown to set nobody as the user owner and the name of the group as the group owner. Then run the **chmod 2770** command on the shared directory.

Pluggable Authentication Modules

? Red Hat Enterprise Linux uses the Pluggable Authentication Modules (PAM) system to check for authorized users.

? PAM modules are called by configuration files in the /etc/pam.d directory. These configuration files are usually named after the service or command that they control.

? There are four types of PAM modules: authentication, account, password, and session management.

? PAM configuration files include lines that list the module_type, the control_ flag, and the path to the actual module, followed by arguments.

? PAM modules are well documented in the /usr/share/doc/pam-*versionnumber*/ txts directory.

Network Authentication Configuration: NIS and LDAP

? NIS allows you to configure one centrally managed username and password database with other Linux and Unix systems on your LAN.

? LDAP provides similar support to NIS, and it supports various forms of encryption.

PREV NEXT

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

User Account Management

1. What's the standard minimum user ID number for regular users on Red Hat distributions?

____

2. What command at a GUI-based text console starts the Red Hat User Manager?

____

Answers

1. The minimum user ID number for regular users on Red Hat distributions is 500. It's 100 on many other Linux distributions.

2. The command in a GUI-based text console that starts the Red Hat User Manager is **system-config-users**.

The Basic User Environment

3. If you want to add files to every new user account, what directory should you use?

____

Answers

3. The directory of files that is automatically added to every new user account is /etc/skel.

Shell Configuration Files

4. The system-wide file associated with the bash shell is _______________.

Answers

4. The system-wide configuration file associated with the bash shell is /etc/bashrc.

Setting Up and Managing Disk Quotas

5. You are running an ISP and provide space for users' Web pages. You want them to use no more than 40MB of space, but you will allow up to 50MB until they can clean up their stuff. How could you use quotas to enforce this policy?

   Set the hard limit to ___________________________

   Set the soft limit to ____________________________

   Enable grace periods with the following command:______________________

6. If you wanted to configure quotas for every user on the /home directory, you'd add the following option to the /home directory line in /etc/fstab:

   -

     ____

Answers

5. Set the hard limit to 40MB; set the soft limit to 50MB. You can set these limits on user1 with the **edquota -u** command; once created, you can apply them to all users with the following command:

```
#
# edquota -p user1 `awk -F: '$3 > 499 {print $1}' /etc/passwd`
```

6. If you wanted to configure quotas for every user on the /home directory, you'd add the **usrquota** option to the /home directory partition directive in /etc/fstab.

Creating and Maintaining Special Groups

7. What command would set the SGID bit on the /home/developer directory?

   ____

8. When creating a special group, should you use the default Group ID for a regular user?

   ____

Answers

7. There are two different commands available to set the SGID bit on the /home/developer directory. If that's all you want to do, run the following command:

```
#
# chmod g+s /home/developer
```

8. When creating a special group, you should not use the default Group ID for a regular user, unless you want to allow others in the group access to all files of that user.

Pluggable Authentication Modules

9. What are the four basic Pluggable Authentication Modules types?

____

____

____

____

10. You are editing the PAM configuration file by adding a module. Which control flag immediately terminates the authentication process if the module succeeds?

____

Answers

[9.](#) The four basic Pluggable Authentication Modules types are: **auth**, **account**, **password**, and **session**.

[10.](#) You are editing the PAM configuration file by adding a module. The **sufficient** control flag immediately terminates the authentication process if the module succeeds.

Network Authentication Configuration: NIS and LDAP

11. What two pieces of information do you need to connect to an NIS server?

____

____

12. If your domain is example.org, what is your LDAP Search Base DN?

dc=_____________________

dc=_____________________

Answers

[11.](#) The two pieces of information you need to connect to an NIS server are the NIS domain name and NIS server (IP address or FQDN).

[12.](#) If your LDAP domain is example.org, your LDAP Search Base DN is **dc=example, dc=org**.

Lab Questions

Lab 1

1. In this first lab, you'll look at setting up automatic connections to a shared network directory. While this lab uses files described in Chapter 10, it is focused on shell configuration files. For the purpose of this lab, assume your username is vaclav and you're mounting a shared NFS /inst directory from a remote computer with an IP address of 192.168.30.4. You're going to mount it in vaclav's home directory, in a blank directory named inst.
    1. Select the regular user of your choice. That user should have files such as .bashrc and .bash_logout.
    2. Find a shared directory on a remote computer.
    3. Set up a local directory for that user as a mount point.
    4. Configure commands for that user to mount and **umount** that remote directory. Make sure those commands run only when that user logs into his or her account.

Answers

1. This lab has two purposes: it is designed to help you understand mounted network directories and the login process. You can substitute the user, the shared network directory, and directories of your choice. To some extent, the automounter described in Chapter 4 provides an alternative. If you have problems, more information on NFS is available in Chapter 10. But based on the premises in this lab, I would take the following steps:
    1. Log in as user vaclav. Create the specified directory. For this lab, you would use the **mkdir/home/vaclav/inst** command.
    2. Test the network connection. Mount the remote NFS directory on the directory that you just created. For this lab, use the following command (substitute the appropriate IP address or host name for your network):

Lab 2

2. In this lab, you will test the quotas created in this chapter. You'll use the basic quota settings described in this chapter and then copy files to fill up the home directory of a user who has a quota applied. The steps required for this lab are straightforward. **?**

1. Set up quotas on the local computer. Use the criteria described earlier in this chapter. If you don't have a separate /home directory partition, you can set up quotas on the top-level root directory (/).
2. Once you've set quotas in your /etc/fstab configuration file, remember to remount the partition where you've created a quota. Alternatively, you could reboot Linux, but that would take time that you may not be able to spare during either of the Red Hat exams.
3. Set up a quota for the user of your choice. Remember that when you use the edquota command on a specific user, you can edit the quota file directly using vi editor commands. Configure a hard and a soft limit for that user.
4. Log in as the user with the quota. Copy some large files to the home directory of that user. To speed up this process, if you want to create an arbitrary large file, say 179MB, run the following command:

Answers

2. The purpose of this lab is to practice creating quotas for users. It's quite possible that you'll have to configure quotas on the Red Hat exams. While you may not have to test quotas in the way described in this lab, it will help you become familiar with the error messages that you'll see when you exceed a hard and then a soft quota limit.

Lab 3

3. In this lab, you'll create a private directory for a group of engineers designing some galleys. You'll want to create a group named galley for the engineers named mike, rick, terri, and maryam. They'll want to share files in the /home/galley directory. What do you need to do? **?**

Answers

3. This is a straightforward process, using the following basic steps:

1. Create accounts for mike, rick, terri, and maryam if required. You can use the **useradd** command, edit the /etc/passwd file directly, or work through the Red Hat User Manager.
2. Set up a group for these users. Configure a group ID outside the range of your regular users with a line such as:

Lab 4

4. You want to make sure even the root user has to enter the root password when opening Red Hat administrative tools. You can do this by modifying the appropriate file in the /etc/pam.d directory. ?

Answers

4. To make the lab work, first review the various system-config-* files in the /etc/pam.d directory. Most (except system-config-lvm and system-config-selinux) include the following three directives:

```
a
auth          include       config-util
a
account       include       config-util
```

PREV NEXT

Chapter 7: System Administration Tools

Overview

Red Hat Enterprise Linux 5 includes the network configuration tools that have made Linux the operating system backbone of the Internet. To configure your network and troubleshoot any network problems you might encounter during the exam, you need the tools described in this chapter.

Printing is a fundamental service for all operating systems. The default print server for RHEL is CUPS, which has replaced the Line Print Daemon. CUPS supports autoconfiguration of shared network printers and includes a Web-based configuration tool. Red Hat also has a customized graphical configuration tool.

Both support connections to printers using other network protocol suites. A lot of system administration is repetitive. Some of it happens when you want to have a "life," more when you'd rather be asleep. In this chapter, you'll learn how to schedule both one-time and periodic execution of jobs. When troubleshooting, system logging often provides the clues that you need to solve a lot of problems.

You may notice that I diverge slightly from the RH300 course outline in Chapters 5 and 7. To me, it made more sense to include Kickstart with the other package management tools described in Chapter 5.

Inside the Exam

This chapter directly addresses four items in the Red Hat Exam Prep guide. While the focus is on RHCT requirements, all RHCE candidates need to remember that their exam includes these requirements. The first Red Hat Exam Prep item is associated with the following RHCT Troubleshooting and System Maintenance skill:

- Diagnose and correct misconfigured networking.

This chapter also addresses the following RHCT Installation and Configuration skills:

- Configure the scheduling of tasks using cron and at.
- Use scripting to automate system maintenance tasks.
- Configure printing.

The final part of this chapter is important for the Troubleshooting and System Maintenance portion of each exam, as log files can help you find the cause of many problems, especially those related to the boot process as well as many services.

If you need to configure SELinux, as defined for the RHCE exam, just about all you can do for the services in this chapter is disable SELinux protection for part or all of each service. However, the SELinux requirements as defined in the Red Hat Exam Prep guide are focused on networking services, which are not covered in this chapter.

Certification Objective 7.01-Network Configuration

The network is where the power of Red Hat Enterprise Linux really comes alive; however, getting there may not be trivial. As in all other things Linux, it's a learning experience. Many critical network configuration settings are stored in the /etc/sysconfig directory.

In most cases, you'll configure networking when you install RHEL during either exam. However, you may encounter and need to diagnose networking problems, especially during the Troubleshooting portion of either exam.

Exam Watch

Learn the configuration files in the /etc/sysconfig/network-scripts/ and /etc/sysconfig/ directories. These are crucial to the configuration of Red Hat Enterprise Linux. If you have a configuration to change or repair, it may involve files in one of these directories. If you have a problem on the troubleshooting exam, you may find the solution in these files. Even if there's an existing well-known configuration file such as httpd.conf, you can find additional configuration options in /etc/sysconfig/httpd. Red Hat has consolidated a number of key configuration files in these directories, so expect them to become even more important in the future.

The configuration file that provides the foundation for others in RHEL 5 networking is /etc/sysconfig/network. It can contain up to six directives, as described in Table 7-1. If you don't see the directives in your /etc/sysconfig/ network file, the situation does not apply. For example, if you don't see the **GATEWAYDEV** directive, you probably have only one network card on your computer.

Table 7-1: /etc/sysconfig/network Variables

Variable	Description
NETWORKING	Can be yes or no, to configure or not configure networking.
NETWORKING_IPV6	Can be yes or no, to configure networking under IPv6.
NISDOMAIN	If you're connected to an NIS network, this should be set to the name of the NIS domain.
HOSTNAME	Sets the host name of the local computer. If you don't see this directive, it may be set by a DHCP server.
GATEWAY	Sets the IP address for the gateway for your network. If you don't see this directive, it may be set by a DHCP server.
GATEWAYDEV	Sets the network device, such as eth0, that this computer uses to reach a gateway. You won't see this if you have only one network card on your computer.

In most cases, /etc/sysconfig/network contains three directives:

```
NETWORKING=yes
NETWORKING_IPV6=yes
```

Certification Objective 7.02-The CUPS Printing System

RHEL comes with one print service, the Common Unix Printing System (CUPS). It's the successor to the Line Print Daemon (LPD) and the companion Line Printer Next Generation (LPRng), which is no longer offered with RHEL or Fedora Linux. You can configure printers directly through the CUPS configuration files in the /etc/cups directory. Alternatively, RHEL includes two quality GUI tools that you can use to configure local and remote printers on your network. One is a Web-based interface; Red Hat has a companion printer configuration utility.

Exam Watch

It doesn't matter how you configure a print server for the Red Hat exams. Whether you use Red Hat's utility or the Web-based tool, or you edit the configuration files directly, all that matters is that you get the result specified on your exam.

CUPS is the Linux/Unix implementation of the Internet Printing Protocol (IPP). I expect IPP to become a fairly universal standard for printer configuration sometime in the future. Microsoft Print servers can now use IPP; Apple systems starting with OS 10.2 use CUPS with IPP.

Installing and Starting CUPS

CUPS and a number of print databases are installed with the Printing Support package group. It includes 10 RPM packages, some of which appear unrelated. If you haven't already installed this package group during the RHEL installation process, it's likely most efficient to install it using the **yum groupinstall printing** command or the Package Manager described in Chapter 5. If you want to learn to install the packages on your own, you can review the packages and groups as listed in the Package Manager.

It's easy to start and configure CUPS to launch when Linux boots on your computer. The **cups** service script works like most of the other services on RHEL. In other words, you can start it with the following command:

```
#
# service cups start
```

Certification Objective 7.03-Automating System Administration: cron and At

The cron system is essentially a smart alarm clock. When the alarm sounds, Linux runs the commands of your choice automatically. You can set the alarm clock to run at all sorts of regular time intervals. Alternatively, the at system allows you to run the command of your choice, once, at a specified time in the future.

RHEL installs the cron daemon (**crond**) by default. It's configured to check the /var/spool/ cron directory for jobs by user. It also checks for scheduled jobs for the computer under /etc/ crontab and in the /etc/cron.d directory.

Exam Watch

Because cron always checks for changes, you do not have to restart cron every time you make a change.

The behavior of the Linux cron is different from Unix, where the cron daemon wakes up only when it needs to launch a program.

The System crontab and Components

The crontab file is set up in a specific format. Each line can be blank, a comment (which begins with #), a variable, or a command. Naturally, blank lines and comments are ignored.

When you run a regular command, the actions of the shell are based on environmental variables. To see the environmental variables, run the **env** command. Some of the standard variables in RHEL include **HOME** as your home directory, **SHELL** as the default shell, and **LOGNAME** as the username.

You can set different variables within the crontab file, or you can set environmental variables with the following syntax:

```
Variable=Value
```

Certification Objective 7.04-Understanding, Maintaining, and Monitoring System Logs

An important part of maintaining a secure system is keeping track of the activities that take place on the system. If you know what usually happens, such as understanding when users log into your system, you can use log files to spot unusual activity. Red Hat Enterprise Linux comes with several utilities you can use to monitor activity on a system. These utilities can help you identify the culprit if there is a problem.

RHEL comes with two logging daemons. The kernel log daemon service, **klogd**, logs kernel messages and events. The **syslogd** daemon logs all other process activity. You can use the log files that **syslogd** generates to track activities on your system. If you are managing multiple Red Hat Enterprise Linux systems, you can configure the **syslogd** daemon on each system to log messages to a central host system.

Both daemons are typically active by default, and both can be activated by the /etc/init.d/syslog script. Once these daemons start, the **syslogd** daemon examines /etc/syslog.conf to find the logging options that you may have configured.

On the Job

The only available SELinux options associated with logging services disable protection; they're associated with the *klogd_disable_trans* and *syslogd_disable_trans* booleans.

System Log Configuration File

You can configure what **syslogd** records through the /etc/syslog.conf configuration file. As shown in Figure 7-15, it includes a set of rules for different facilities (if the corresponding packages are installed): authpriv, cron, kern, mail, news, user, and uucp. You may not see everything that's shown in Figure 7-15; for example, if you haven't installed the Internet News Network (INN) service, you won't see the associated directives in this file.

```
# Log all kernel messages to the console.
# Logging much else clutters up the screen.
#kern.*                                                  /dev/console

# Log anything (except mail) of level info or higher.
# Don't log private authentication messages!
*.info;mail.none;news.none;authpriv.none;cron.none          /var/log/message
s

# The authpriv file has restricted access.
authpriv.*                                               /var/log/secure

# Log all the mail messages in one place.
mail.*                                                   -/var/log/maillog


# Log cron stuff
cron.*                                                   /var/log/cron

# Everybody gets emergency messages
*.emerg                                                  *

# Save news errors of level crit and higher in a special file.
uucp,news.crit                                           /var/log/spooler

# Save boot messages also to boot.log
local7.*                                                 /var/log/boot.log

#
# INN
#
news.=crit                                         /var/log/news/news.crit
news.=err                                          /var/log/news/news.err
news.notice                                        /var/log/news/news.notice
```

Figure 7-15: The syslog.conf log configuration file

Each facility is associated with several different levels of logging, known as the *priority*. In ascending order, log priorities are **debug**, **info**, **notice**, **warn**, **err**, **crit**, **alert**, **emerg**. There's also a generic **none** priority that logs all messages at all levels.

For each facility and priority, log information is sent to a specific log file. For example, consider the following line from /etc/syslog.conf:

```
*
*.info;mail.none;news.none;authpriv.none;cron.none /var/log/messages
```

Certification Summary

You've read about the configuration files in the /etc/sysconfig hierarchy. Some are important for networking, and others are important for basic parameters such as the system clock, mouse, and keyboard. You also learned about a number of related networking commands, including **ifup**, **ifdown**, ifconfig, netstat, **arp**, and **dhclient**.

A number of important network client services are associated with Red Hat Enterprise Linux. CUPS supports the configuration of printers locally or over the network. With CUPS, the configuration files are stored primarily in the /etc/cups directory. However, it also includes a list of printers in /etc/printcap to accommodate sharing through Samba. You can edit the CUPS configuration files directly with Red Hat's Printer Configuration tool or the CUPS Web-based interface.

The cron and at daemons can help you manage when and how services start on your system. Finally, log files can be configured to collect data from any number of services.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 7.

Network Configuration

- ? Key network configuration files are in the /etc/sysconfig directory: the *network* file, and the networking and network-scripts subdirectories.
- ? You can start the Network Configuration utility with the **system-config-network** command.
- ? To manage network settings on each interface, use **dhclient**, **ifup**, and **ifdown**.
- ? The ifconfig command is used to configure and display network devices.
- ? Use **ifup eth0** and **ifdown eth0** to activate and deactivate the eth0 interface.
- ? The netstat command is used to display a plethora of network connectivity information; **route -n** is another way to check the current routing table.
- ? The **arp** command is used to view or modify the local hardware address database.

The CUPS Printing System

- ? The Printer Configuration tool, which you can start with the **system-config-printer** command, can be used to configure most popular printers in /etc/cups/cupsd.conf.
- ? CUPS provides a Web-based interface. Once enabled, you can get to this interface in your browser by navigating to http://localhost:631.

Automating System Administration: cron and At

- ? The cron system allows any user to schedule jobs so they run at given intervals.
- ? The at system allows users to configure jobs to run once at a scheduled time.
- ? The crontab command is used to work with cron files. Use **crontab -e** to edit, **crontab -l** to list, or **crontab -d** to delete cron files.

? The /etc/cron.allow and /etc/cron.deny files are used to control access to the cron job scheduler.

Understanding, Maintaining, and Monitoring System Logs

? Red Hat Enterprise Linux includes two logging daemons: **klogd** for kernel messages and **syslogd** for all other process activity. Both are activated by the /etc/init.d/syslog service script.

? You can use log files generated by the **syslogd** daemon to track activities on your system.

? Most log files are stored in /var/log.

? You can configure what is logged through the syslog configuration file, /etc/ syslog.conf.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams.

Network Configuration

1. What three directives are typically found in the /etc/sysconfig/network file? ?

   ____

   ____

   ____

2. What command would you use to assign an IP address of 192.168.99.44 to your eth1 network card? ?

   ____

3. What command returns the current routing table? ?

   ____

4. What command deactivates the eth0 network device? ?

   ____

The CUPS Printing System

5. You want to look at your current printer configuration in a GUI desktop interface. What starts the Red Hat GUI Printer Configuration tool? ?

   ____

6. You want to configure a group of printers as a printer class. Which GUI tool can you use for this purpose? ?

   ____

7. What is the main CUPS printer configuration file? ?

   ____

8. What command can you use to remove print job number 12 from the print queue? ?

   ____

Answers

5. The command that starts the Red Hat GUI Printer Configuration tool from a GUI command line is **system-config-printer**.

6. To configure a group of printers as a printer class, you can use the CUPS Web-based configuration tool.

7. The main CUPS printer configuration file is /etc/cups/cupsd.conf.

8. The command that removes print job number 12 from the default print queue is

```
#
# lprm 12
```

Automating System Administration: cron and At

9. You want to schedule a maintenance job, maintenance.pl, to run from your home directory on the first of every month at 4:00 A.M. You've run the **crontab -e** command to open your personal job file. Assume you've added appropriate PATH and **SHELL** directives. What directive would you add to run the specified job at the specified time?

____

10. If you see the following entry in your /etc/crontab:

```
4
42 4 1 * * root run-parts /etc/cron.monthly
```

Answers

9. The directive that runs the maintenance.pl script from a home directory at the noted time is

```
0
00 4 1 * * ~/maintenance.pl
```

10. Based on the noted entry in /etc/crontab, the next time Linux will run the jobs in the /etc/cron .monthly directory is on the first of the upcoming month, at 4:42 A.M.

Understanding, Maintaining, and Monitoring System Logs

11. Assume you normally work from a user account called sysadm. What entry might you add to /etc/syslog.conf to notify you whenever there is a serious problem with the kernel?

____

12. What three log files in /var/log are most relevant to the installation process?

____

____

____

Answers

11. The entry that you might activate in /etc/syslog.conf to notify you whenever a serious problem with the kernel occurs is

```
k
kern.* /dev/console
```

12. The log files in /var/log that are most relevant to the installation process are **anaconda.log**, **anaconda.syslog**, and **anaconda.xlog**. If you install in text mode, anaconda.xlog won't be created.

Lab Questions

Lab 1

1.

1. ?

In this lab, you'll deliberately misconfigure your network card, observe the results, and see what you can do to restore the original configuration. I assume that you have a computer on a network, as that is a basic requirement of the Red Hat exams. If you have more than one network card, make sure you select the card that you use to connect to outside networks.

This lab assumes a basic understanding of IPv4 network addresses. It also assumes that you have a DHCP server on your system. (For now, it doesn't matter what operating system has the DHCP server; many home routers serve this purpose.) To prepare this lab, take the following steps:

1. Boot your system into a standard runlevel with networking, 3 or 5.
2. Inspect the current configuration of your network card with the ifconfig command. Inspect the current routing table, as specified by the **route -n** (or **netstat -nr**) command. If you have several network devices, find the device associated with your gateway.
3. Select the network device that connects to external networks, such as eth0. Change the IP address to a private IPv4 address on a different network. If you have multiple network cards, make sure the new address does not correspond to any existing network.
4. Confirm the change by running ifconfig again.
5. Try the **ping** command to the new IP address.
6. Try the **ping** command to an IP address on an external computer. It doesn't matter if the external computer is on a local network or the Internet. What message do you see?
7. Run the **ping** command to a host name of an external computer. It doesn't matter if the external computer is on a local network or the Internet. What message do you see?
8. Run the **route -n** command again. Is there a default route?
9. Assuming the network device you've changed is eth0, run the **dhclient eth0** command. What do you see? Run the ifconfig command again. Has the IP address changed?
10. Search through active processes for the **dhclient** command. What did you find?
11. Kill any active **dhclient** command process.
12. Apply the **dhclient** command to the target network device.
13. Redo step 2. What happened to the routing table? What about your network interface?

Answers

1. 1.

This set of "answers" explain what you might expect during each step.

1. The standard runlevel for console logins is 3; the standard runlevel for GUI logins is 5. To boot into the runlevel of your choice, press a key when you see "Booting Red Hat Enterprise Linux 5 ?": press P and enter the GRUB password if required, and press A to modify the kernel arguments. (This works only on a non-Xen kernel option.) You can then specify the boot runlevel by adding it to the end of the kernel command line.

2. If you have several network devices, you can focus by specifying a device with a command such as **ifconfig eth0**. The **route -n** command should collate connected networks with the appropriate network device (**netstat -nr** works just like **route -n**).

3. Make sure the IP address you select is on a different network. It's best if you use private IP addresses. Some address ranges are listed in Chapter 1.

4. When you run ifconfig again, you'll see the change you made.

5. When you try the **ping** command to the new IP address, it should work.

6. When you try the **ping** command to an IP address on an external computer, you'll see an error such as:

Lab 2

2. In this lab, you'll want to use the Red Hat Printer Configuration utility to connect a printer to your Linux system. The printer can be local or remotely connected through your LAN. As you'll want to create a printer class, you'll need more than one printer. (Hint: You can configure the same printer as many times as necessary, as long as you use different names.)

If you have only one physical printer, you can set up multiple print queues with different printer names. CUPS sees each print queue as if it were a separate printer.

Once you've created multiple printers, return to the Printer Configuration tool. Run the **system-config-printer** command in a GUI. Navigate to http://localhost:631. Click Add Class. Follow the prompts to create a printer class with the printers that you've configured.

Once you've created a new printer class, inspect the result in the /etc/cups/classes.conf configuration file. Check the contents of your /etc/printcap and /etc/printers.conf files. What are the names of the printers that you see? Do you see any surprises in the list?

If you have a Microsoft Windows computer on your network, activate the Samba service if required. Check the printer names as shown in the browse list. Do you see any surprises on this list? (For more information on Samba, see Chapter 10.)

Answers

2. Starting the Printer Configuration utility is easy. One way in the GUI is to press ALT-F2 and then type **system-config-printer** in the text box that appears. Then you can click the New Printer button to start a configuration wizard. If you like, you can create different print queues for the same printer. Just repeat the same process, using a different printer name.

When you've created two or more printers, click New Class. This starts a similar configuration wizard, where you can collect preconfigured printers of your choice in a print class. As with new printers, a print class requires a name, description, and location.

Click Forward; select more than one computer as members of the printer class. Click Forward; once your new printer class is confirmed, click Apply.

You should now see the printer class and member printers in the /etc/cups/classes.conf file. You'll find a list of printers in /etc/printcap; you'll find the names of any configured printers *and* printer classes in this file. You should also see the list of printers in /etc/printcap in any Microsoft Windows Network Neighborhoods or My Network Places that is connected to the same network. This assumes that you've activated a Samba server on the local print server computer, of course.

Lab 3

3. As the root user, create jobs that change the login message for users at the text console. To do so, you'll want to change the content of /etc/motd. Make sure that people who log in at different times get appropriate messages:

- If users log in between 7 A.M. and 1 P.M., create the login message "Coffee time!"
- If users log in between 1 P.M. and 6 P.M., create the login message "Want some ice cream?"
- If users log in between 6 P.M. and 7 A.M., create the login message "Shouldn't you be doing something else?"

Answers

3. To modify the login messages as noted, take the following steps:

1. Log in as the root user.
2. Run the **crontab -e** command.
3. Add appropriate environment variables, at least the following:

Lab 4

4. In this lab, you'll find the value of several different log files. In preparation, use the wrong password to log into a regular account. Then take the following steps: ?

1. Navigate to /var/log as the root user.
2. Explore the contents of the anaconda.* log files.
3. Run the **utmpdump btmp** command. Do you see the login attempt? Can you tell if it succeeded?
4. Review the contents of the cron log file. Scroll through it. If your computer has been on for a while, most of what you see will be based on the **run-parts /etc/cron.hourly** command. Alternatively, if you reboot on occasion, you'll see messages associated with the anacron service.
5. Review the contents of the dmesg log file. Compare the beginning of it with the start of the anaconda.syslog file. Which one includes the currently booted kernel?
6. Navigate toward the bottom of the dmesg file. Can you identify the amount of swap space? What about one or more partitions with the default EXT3 filesystem?
7. Review the maillog log file. Do you see any messages associated with mail messages? If there are a lot of messages associated with the root account, run the **mail** command (to exit from the mail prompt [**&**], press CTRL-D).
8. Review the secure log file. Navigate to the bottom of the file. Do you see a message associated with the failed login?
9. Finally, review the Xorg.0.log file. Do you see any messages related to the mouse (or other pointing device) near the end of the file? How does that work when you didn't configure a pointing device during the installation process?

Answers

4. There are no secret solutions in this lab; the intent is to get you to review the contents of key log files to see what should be there.

When you review the anaconda.* files in /var/log and compare them to other files, you may gain some insight on how to diagnose installation problems. In future chapters, you'll examine some of the log files associated with specific services; many are located in subdirectories such as /var/log/samba/ and /var/log/httpd/.

The failed login should be readily apparent in the /var/log/secure file. You may be able to get hints in the output to the **utmpdump btmp** command.

When you review the /var/log/cron file, you'll see when standard cron jobs were run. Most of the file should be filled (by default) by the standard hourly job, **run-parts /etc/cron.hourly**, from the /etc/crontab configuration file. If you've rebooted, you may see the anacron service, and you should be able to search for the job of the same name.

While /var/log/dmesg includes the currently booted kernel, it may be the same kernel as the one associated with /var/log/anaconda.syslog, if you haven't upgraded kernels. At the end of /var/log/dmesg, you can find the filesystems mounted to the EXT3 format, as well as currently mounted swap partitions. For example, the following lists swap partitions that happen to be on two different hard drives:

```
EXT3 FS on sda9, internal journal
EXT3-fs: mounted filesystem with ordered data mode.
Adding 97995k swap on /dev/sda3. \
 Priority:-1 extents:1 across:97995k
```

1. The three directives typically found in the /etc/sysconfig/network file are **NETWORKING**, **NETWORKING_IPV6**, and **HOSTNAME**. If you allow a DHCP server to assign host names, the HOSTNAME directive isn't required.

2. The command that assigns an IP address of 192.168.99.44 to an eth1 network card is

Chapter 8: Kernel Services and Configuration

Overview

In this chapter, you'll learn how to upgrade standard kernels as well as configure, compile, and install your own custom kernels. You'll see several different ways to customize and optimize your kernel configuration for size and functionality. Finally, you'll examine recommended techniques for configuring and installing the kernel.

You'll also discover how to manage and modify special partitions associated with RAID arrays and LVM filesystems. While it's most efficient to configure these partitions during the installation process, you may have to modify them during your exam.

In several places in this chapter, I embed a command such as **'uname -r'** in the name of a directory or file. This command substitutes itself in the name of directory or file. If unsure, run it in your own system.

Inside the Exam

Managing Kernels

As a competent Linux administrator, you need to know how to install, patch, and recompile kernels. It's easy to install a new kernel from an RPM, which makes it a reasonable requirement on the RHCT and RHCE exams. Early versions of the RHCE Exam Prep guide suggested that you need to know how to recompile the Linux kernel. But that was a long process! However, as the process is now much easier, I would not be surprised to see the requirement return for RHCEs in the near future.

In addition, the current Exam Prep guide suggests that you need to know how to do the following during the Installation and Configuration portion of both exams:

- Properly update the Kernel package.
- Use /proc/sys and sysctl to modify and set kernel run-time parameters.

RAID and LVM

The Exam Prep guide also describes skills associated with configuring RAID and LVM after installation. I've included it in this chapter to match the RH300 course outline more closely. While it's easiest (and therefore best) if you can configure RAID and LVM during the installation process, it's not always possible. If you make a mistake during the installation process, you don't need to start over. The Exam Prep guide suggests that during the Troubleshooting and System Maintenance portion of the exam, RHCTs need to know how to

- Add new partitions, filesystems, and swap to existing systems.

RHCEs need to know how to

-

Add, remove, and resize logical volumes. Remember, if you're taking the RHCE exam, you also need to meet all RHCT requirements.

Certification Objective 8.01-The Basics of the Kernel

The kernel is the heart of the operating system. It manages communication with hardware, decides which processes to run, and provides each process with an isolated, virtual address space in which to run. The kernel is what the GRUB boot loader loads into memory. The kernel loads device driver modules. It also allocates hardware resources such as IRQ ports, I/O addresses, and DMA channels. A recompiled kernel can lead to:

- Greatly improved speed at which kernel services operate.
- Direct support for commonly used drivers.
- Dynamic loading of appropriate drivers as modules.
- Lower memory consumption by removing unneeded components.
- Support for high-end hardware, such as memory above 4GB, hardware array controllers, symmetric multiprocessing (multiple CPU) support, and more.

In essence, you can customize the Linux kernel any way you want. The best way to do it is to make it fit every detail of installed hardware. However, you may not need to be so picky. In many cases, all you need to do is install the updated kernel RPM. In other cases, such as compiling third-party drivers, all you need to install is the corresponding kernel-devel RPM.

Exam Watch

Xen is based on a specially customized Linux kernel for virtual machines. The files associated with the Xen-based kernel are different from regular kernels. While Xen is important to RHEL 5, I don't believe you'll have to install or use it for your systems on the RHCE exams, at least until RHEL 6 is released. Of course, I could be wrong. Monitor the Red Hat Exam Prep guide for the latest information.

Best Practices

You should compile your kernel with only the elements you need. The more that is left out, the faster the whole system will run. For example, if there is no sound card, sound card support can be removed from the kernel. By removing unneeded devices, you will:

- Decrease the size of the kernel.
- Provide a modest increase in speed for the devices that are present.
-

Make more hardware resources available for other hardware such as network cards, disk controllers, and so on.

- Reduce the risk of hardware limits, such as those that may be based on the size of the compressed kernel.

But don't remove things you don't understand, as those components may be essential to the smooth functioning of the kernel.

Generally, it is a good idea to have device drivers compiled as modules for any equipment that you may add in the near future. For example, if you may use your Linux computer as a router, you'll need a second network card, and you can add support for that card to your kernel. For example, if you have a 3Com 3c595 network card installed but you also have some 3Com 3c905 cards in storage, it may be a good idea to include the 3c905 module. That way, you can simply swap in the new card and let the module load, causing minimum downtime.

Modules are kernel extensions. They are not compiled directly into the kernel but can be plugged in and removed as needed. When configured as a module, a hardware failure such as that of a network card will not cause the whole system to fail.

Kernel Concepts

You will need to understand some basic kernel concepts before you can compile your own kernel. Kernels can be organized as one big unit or as a lot of interconnected pieces. Kernels are called up by boot loaders when you start your system.

Monolithic Versus Modular

A *monolithic* kernel is a kernel in which all the device modules are built directly into the kernel. *Modular* kernels have many of their devices built as separate loadable modules. Monolithic kernels can communicate with devices faster, since the kernels can talk to the hardware only indirectly through a module table. Unfortunately, the typical monolithic kernel is huge, which reduces available RAM. In addition, some systems just can't boot a kernel that's too large.

Linux once had problems loading modular kernels for some hardware. With a monolithic kernel, the drivers are already there and are often more appropriate for certain components such as embedded hardware.

A modular kernel has greater flexibility. You can compile almost all drivers as modules, and then each module can be inserted into the kernel whenever you need it. Modules keep the initial kernel size low, which decreases the boot time and improves overall performance. If Linux has trouble loading a kernel module, you can use the modprobe or **insmod** command to load modules as needed, and add those options to the /etc/modprobe.conf file.

Updating the Kernel

Updating the kernel is not as difficult as it looks. You should never overwrite or upgrade an existing kernel, as mistakes happen. New kernels are handled by installing the newly built kernel in /boot and then adding another boot option to your boot loader configuration file (/boot/grub/grub.conf) for the new kernel. GRUB treats the new kernel as if it were an entirely new operating system.

If you install the new kernel directly from a Red Hat configured RPM, it updates the boot loader automatically. Chapter 5 explored this process briefly.

If you do make a drastic mistake and the kernel doesn't boot, you can simply reboot the system and select the old kernel from the GRUB menu. You should also save existing kernel configuration files so that you have a template for newer kernels. This is discussed in more detail later in this chapter.

Other RHEL Kernels

There are a number of different kernels included with the RHEL installation files. You can and should install the kernel best suited to your system. Available RHEL 5 kernels are briefly discussed in Table 8-1. For the real *versionnum*, run the **uname -r** command. To verify your *arch*, or architecture (such as i686), run the **uname -m** command. As described in the table, there are different versions of kernel-devel, kernel-PAE, kernel-xen, and kernel-headers packages for each supported architecture.

Table 8-1: Available Red Hat Enterprise Linux Kernels (and Related Packages)

Kernel RPM	**Description / Architecture**
kernel-*versionnum*.i686	Designed for PCs with a single Intel/AMD CPU; also works with dual-core systems
kernel-*versionnum*.ia64	Designed for Itanium2 systems
kernel-devel-*versionnum*	Installs drivers and other information to help compile third-party drivers
kernel-PAE-*versionnum*	If you have more than 4GB of RAM, install the PAE kernel associated with your CPU architecture
kernel-PAE-devel-*versionnum*	If you have more than 4GB of RAM, install the PAE kernel associated with your CPU architecture
kernel-headers-*versionnum*	Includes kernel headers; often sufficient for drivers
kernel-*versionnum*.src.rpm	Includes the source code for the RHEL kernel

I don't list all available RHEL architectures in Table 8-1, and list them just for the basic kernel packages. Remember, the focus of the Red Hat exams is still based on the basic 32-bit Intel/AMD/clone CPU. PPC and s/390 systems are not (yet) supported in RHEL 5.

The table provides just a short list of kernel packages available for RHEL 5. It does not include Xen-related kernels. For more information on RHEL kernels available for multi-CPU or higher-end CPUs, refer to the RHEL documentation available online from www.redhat.com/docs/manuals/enterprise/.

The /boot Partition

The Linux kernel is stored in the partition with the /boot directory. New kernels must also be transferred to this directory. By default, RHEL configures a partition of about 100MB for this directory. This provides enough room for your current kernel plus several additional upgraded kernels.

The /proc Filesystem

The /proc directory is based on a virtual filesystem; in other words, it does not include any files that are stored on the hard drive. But it is a window into what the kernel sees of your computer. It's a good idea to study the files and directories in /proc, as it can help you diagnose a wide range of problems. Figure 8-1 shows the /proc directory from a typical RHEL computer.

```
[michael@enterprise5dl ~]$ \ls /proc/
1     1954  2364  2835  2927  3271  5            ide         partitions
10    1960  2365  2838  2931  329   6            interrupts  schedstat
1024  1981  2381  2839  2936  332   67           iomem       scsi
11    2     2382  2846  2949  3451  7            ioports     self
135   2033  2388  2849  2951  3501  70           irq         slabinfo
136   2058  2393  2851  2957  3505  72           kallsyms    stat
137   2082  2405  2865  2962  3506  acpi         kcore       swaps
138   2106  2455  2870  2964  3509  buddyinfo    keys        sys
1618  2122  2459  2872  2966  3553  bus          key-users   sysrq-trigger
1720  2127  2460  2876  2990  3640  cmdline      kmsg        sysvipc
1737  2144  2463  2878  2992  365   cpuinfo      loadavg     tty
1739  2175  2466  2882  2993  3818  crypto       locks       uptime
1758  2217  2467  2885  3     3861  devices      mdstat      version
1761  2226  2486  2899  3020  3865  diskstats    meminfo     vmcore
1798  2243  2515  290   319   3867  dma          misc        vmstat
1821  2259  2688  2906  3239  3869  driver       modules     zoneinfo
1839  2297  2718  2908  324   399   execdomains  mounts
1865  2313  2722  2915  3241  4     fb           mpt
1900  2328  2725  2918  3242  4004  filesystems  mtrr
1937  2347  2776  2923  3270  4008  fs           net
[michael@enterprise5dl ~]$
```

Figure 8-1: A Red Hat Enterprise Linux/proc directory

The numbered items are based on process IDs. For example, the process ID of init is 1. The files in this directory include the memory segments that make up the active process. The contents of each of these files include the active memory for that process.

The other items in the listing are files and directories that correspond to configuration information for components such as DMA channels or whole subsystems such as memory information.

Take a look at some of these files. For example, the /proc/meminfo file provides excellent information as to the state of memory on the local computer, as shown in Figure 8-2. It can help you determine whether RHEL is having trouble detecting all of the memory on your computer.

```
MemTotal:        513432 kB
MemFree:          46768 kB
Buffers:          21332 kB
Cached:          297780 kB
SwapCached:           0 kB
Active:          215492 kB
Inactive:        207516 kB
HighTotal:            0 kB
HighFree:             0 kB
LowTotal:        513432 kB
LowFree:          46768 kB
SwapTotal:      1048568 kB
SwapFree:       1048568 kB
Dirty:             9504 kB
Writeback:            0 kB
AnonPages:       103644 kB
Mapped:           49080 kB
Slab:             31520 kB
PageTables:        3984 kB
NFS_Unstable:         0 kB
Bounce:               0 kB
CommitLimit:    1305284 kB
Committed_AS:    340908 kB
VmallocTotal:    507896 kB
VmallocUsed:       3312 kB
VmallocChunk:    504480 kB
HugePages_Total:      0
HugePages_Free:       0
HugePages_Rsvd:       0
Hugepagesize:      4096 kB
```

Figure 8-2: Detected memory information

It can also help you measure the current memory state of your system. For example, if your system is overloaded, you'll probably find very little free swap space. The **HugePages** settings are associated with systems with over 4GB of RAM.

Now you can examine how Linux looks at your CPU in the /proc/cpuinfo file, as shown in Figure 8-3. In this particular case, the cpu family information is important; the cpu family value of 6 in this figure corresponds to a 686 CPU. If you have a dual-core CPU (and both cores are detected), you'll see two entries, even if you have only one physical CPU.

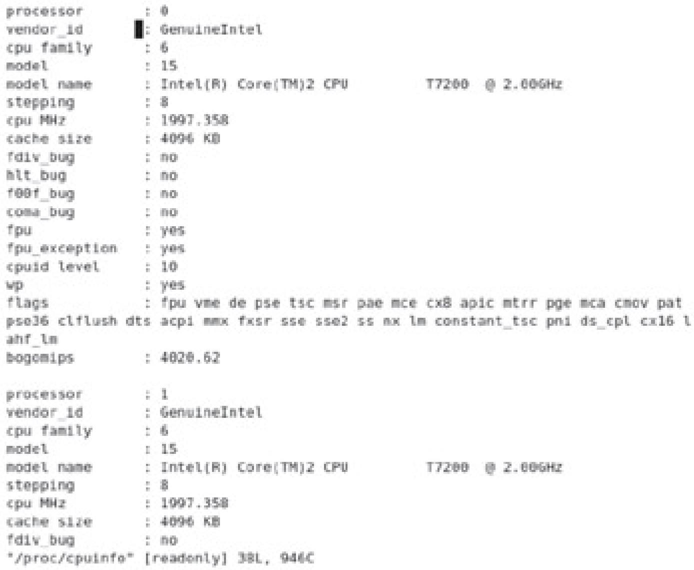

```
processor       : 0
vendor_id       : GenuineIntel
cpu family      : 6
model           : 15
model name      : Intel(R) Core(TM)2 CPU         T7200  @ 2.00GHz
stepping        : 8
cpu MHz         : 1997.358
cache size      : 4096 KB
fdiv_bug        : no
hlt_bug         : no
f00f_bug        : no
coma_bug        : no
fpu             : yes
fpu_exception   : yes
cpuid level     : 10
wp              : yes
flags           : fpu vme de pse tsc msr pae mce cx8 apic mtrr pge mca cmov pat
pse36 clflush dts acpi mmx fxsr sse sse2 ss nx lm constant_tsc pni ds_cpl cx16 l
ahf_lm
bogomips        : 4020.62

processor       : 1
vendor_id       : GenuineIntel
cpu family      : 6
model           : 15
model name      : Intel(R) Core(TM)2 CPU         T7200  @ 2.00GHz
stepping        : 8
cpu MHz         : 1997.358
cache size      : 4096 KB
fdiv_bug        : no
"/proc/cpuinfo" [readonly] 38L, 946C
```

Figure 8-3: Detected CPU information

Many programs are available that simply look at the information stored in /proc and interpret it in a more readable format. The **top** utility is a perfect example. It reads the process table, queries RAM and swap usage and the level of CPU use, and presents it all on one screen.

IP Forwarding

More importantly, there are kernel variables that can be altered to change the way the kernel behaves while it's running. Sometimes it's appropriate to configure a Linux computer as a router between networks. By default, it does not forward TCP/ IP information. You can confirm this with the following command:

```
#
# cat /proc/sys/net/ipv4/ip_forward
 0
```

Certification Objective 8.02-New Kernels, the Easy Way

On the Red Hat exams, you may expected to upgrade your kernel by installing the latest Red Hat kernel RPM. You may be able to patch an existing kernel.

Before we begin, it's important to understand the way kernels are numbered.

Understanding Kernel Version Numbers

The version number associated with the RHEL kernel may look a little confusing, but it tells you a lot about the history of the kernel. The standard RHEL kernel is a version with a number like 2.6.18-8.el5, formatted in a *majorversion.majorrevision .patch-build* format.

The first number (2) is the major version number. These versions provide drastic changes to the kernel. Typically, older version software will *not* work in the newer version when this number changes. Kernel major version numbers are reserved for completely new kernel designs.

The second number (6) actually has two meanings. First, it indicates this is the sixth major revision of major version 2 of the kernel. Second, since it is an even number, it indicates that the kernel release is a stable release. Before version 2.6, an odd second number would indicate a developmental kernel, not suitable for production computers. Now, kernel version 2.7 will also be a production kernel.

On the Job

To promote stability, Red Hat usually works from a slightly older, and presumably more stable, version of the Linux kernel. This is consistent with the demands of the Red Hat customer base; most business customers want to stay away from the "bleeding edge."

The third number (18) is the patch version number for the kernel. These changes are typically small changes, bug fixes, security fixes, and enhancements. Generally, you can use the **zcat** command to increment one patch at a time. For example, if your current kernel is version 2.6.18, you can use the patch-2.6.19.gz file to upgrade your kernel to version 2.6.19.

The fourth number (-8) is a number added by Red Hat. This is the eighth Red Hat version of Linux kernel 2.6.18, which incorporates features customized for Red Hat Enterprise Linux. In some cases, there will be a fifth and even a sixth number that indicates the build number as created by Red Hat. The final bit may be something like el5 or el5xen.

On the Job

Stock kernels use a slightly different four-number system. The first two numbers are identical to the Red Hat system. The third number specifies a major patch; the fourth number specifies a bugfix or security update. If you use a stock kernel, it may overwrite Enterprise-level custom features developed for RHEL 5.

Upgrading Kernels

During the lifetime of any version of RHEL, you may run across a security advisory that strongly suggests that you upgrade your Linux kernel. In this case, a Red Hat kernel RPM will be available through the Red Hat Network.

Exam Watch

You won't have access to the Internet during the Red Hat exams, and therefore, you may not be able to get to the Red Hat Network for updates. However, the Exam Prep guide suggests that you may be required to install an upgraded kernel.

Upgrading a kernel from a Red Hat RPM is fairly easy. Basically, all you need to do is install the new kernel with the appropriate **rpm** or yum command. When properly configured, the RPM automatically upgrades your default boot loader as well. For example, say you've just downloaded the newest kernel RPM from the Red Hat Network to the /tmp directory.

Exam Watch

If you're told to upgrade a new kernel, you'll probably use the *rpm -i* kernel.*rpm* command, and not *rpm -U* kernel.*rpm*. Installing (and not upgrading) newer kernels allows you to boot into the older kernel-in case the new kernel does not work for you.

Be careful. Install (**-i**), don't upgrade (**-U**) your new kernel. Otherwise, if you have a problem, you won't be able to go back to the old working kernel. Installing (**-i**) a new kernel with a command such as

```
#
# rpm -i /tmp/kernel-2.6.18-2.2.1.i686.rpm
```

Certification Objective 8.03-Kernel Sources

One of the strengths of Linux is the ease with which the kernel can be customized to meet your precise needs. But before starting this process, you need the Linux kernel source code. Kernel modules and associated configuration files are covered in this section.

If you choose to recompile the Linux kernel, you'll need several GB of free space available in the partition or volume that contains the /usr directory.

Exam Watch

While the Red Hat Exam Prep guide no longer includes references to recompiling the kernel, you may still need to find kernel modules and configuration files. And I would not be surprised to see a requirement to recompile the kernel return in future RHCE exams.

The Kernel Source Tree and Documentation

When you install the generic kernel source code, it's normally installed in (or linked to) the /usr/src/linux directory. If you install the source code from a Red Hat/Fedora source RPM, the code gets installed in the /usr/src/redhat/BUILD/kernel-2.6.18/ linux-2.6.18.i386/ subdirectory. It can be helpful to link that directory to /usr/src/ linux. Once linked, the /usr/src directory should look similar to the following:

```
#
# ls -l /usr/src/
 t
 total 20
 d
 drwxr-xr-x 3 root root 4096 Mar 14 09:32 kernels
 l
```

PREV NEXT

Certification Objective 8.04-Recompiling a Kernel

While references to recompiling the Linux kernel have been removed from the Red Hat exam requirements, RHCEs in the real world are expected to know how to perform high-level tasks such as optimizing and recompiling the Linux kernel.

This section looks at the kernel configuration file and then proceeds with a discussion of the different tools available to edit the kernel configuration. Finally, you'll see the commands needed to compile your new configuration into the kernel, the files added to the /boot directory, and the settings added to the boot loader.

The Kernel Configuration Scripts

After you've configured a kernel once, the configuration information is stored in a hidden file, .config, in the Linux source code directory. It is structured as a listing of variables. Here are some entries from the .config file:

```
C
CONFIG_NETDEVICES=y
C
CONFIG_BONDING=m
```

Certification Objective 8.05-Advanced Partitioning: Software RAID

A Redundant Array of Independent Disks (RAID) is a series of disks that can save your data even if a catastrophic failure occurs on one of the disks. While some versions of RAID make complete copies of your data, others use the so-called parity bit to allow your computer to rebuild the data on lost disks.

Linux RAID has come a long way. A substantial number of hardware RAID products support Linux, especially those from name-brand PC manufacturers. Dedicated RAID hardware can ensure the integrity of your data even if there is a catastrophic *physical* failure on one of the disks. Alternatively, you can configure software-based RAID on multiple partitions on the same physical disk. While this can protect you from a failure on a specific hard drive sector, it does not protect your data if the entire physical hard drive fails.

Depending on definitions, RAID has nine or ten different levels, which can accommodate different levels of data redundancy. Combinations of these levels are possible. Several levels of software RAID are supported directly by RHEL: levels 0, 1, 5, and 6. Hardware RAID uses a RAID controller connected to an array of several hard disks. A driver must be installed to be able to use the controller. Most RAID is hardware based; when properly configured, the failure of one drive for almost all RAID levels (except RAID 0) does not destroy the data in the array.

Linux, meanwhile, offers a software solution to RAID. Once RAID is configured on a sufficient number of partitions, Linux can use those partitions just as it would any other block device. However, to ensure redundancy, it's up to you in real life to make sure that each partition in a Linux software RAID array is configured on a different physical hard disk.

On the Job

The RAID md device is a meta device. In other words, it is a composite of two or more other devices such as /dev/hda1 and /dev/hdb1 that might be components of a RAID array.

The following are the basic RAID levels supported on RHEL.

RAID 0

This level of RAID makes it faster to read and write to the hard drives. However, RAID 0 provides no data redundancy. It requires at least two hard disks.

Reads and writes to the hard disks are done in parallel-in other words, to two or more hard disks simultaneously. All hard drives in a RAID 0 array are filled equally. But since RAID 0 does not provide data redundancy, a failure of any one of the drives will result in total data loss. RAID 0 is also known as *striping without parity*.

RAID 1

This level of RAID mirrors information between two disks (or two sets of disks-see RAID 10). In other words, the same set of information is written to each disk. If one disk is damaged or removed, all of the data is stored on the other hard disk. The disadvantage of RAID 1 is that data has to be written twice, which can reduce performance. You can come close to maintaining the same level of performance if you also use separate hard disk controllers, which prevents the hard disk controller from becoming a bottleneck. RAID 1 is relatively expensive. To support RAID 1, you need an additional hard disk for every hard disk worth of data. RAID 1 is also known as *disk mirroring*.

RAID 4

While this level of RAID is not directly supported by the current Linux distributions associated with Red Hat, it is still supported by the current Linux kernel. RAID 4 requires three or more disks. As with RAID 0, data reads and writes are done in parallel to all disks. One of the disks maintains the parity information, which can be used to reconstruct the data. Reliability is improved, but since parity information is updated with every write operation, the parity disk can be a bottleneck on the system. RAID 4 is known as *disk striping with parity*.

RAID 5

Like RAID 4, RAID 5 requires three or more disks. Unlike RAID 4, RAID 5 distributes, or *stripes*, parity information evenly across all the disks. If one disk fails, the data can be reconstructed from the parity data on the remaining disks. RAID does not stop; all data is still available even after a single disk failure. RAID 5 is the preferred choice in most cases: the performance is good, data integrity is ensured, and only one disk's worth of space is lost to parity data. RAID 5 is also known as *disk striping with parity*.

RAID 6

RAID 6 literally goes one better than RAID 5. In other words, while it requires four or more disks, it has two levels of parity and can survive the failure of two member disks in the array.

RAID 10

I include RAID 10 solely to illustrate one way you can combine RAID levels. RAID 10 is a combination of RAID 0 and RAID 1, which requires a minimum of four disks. First, two sets of disks are organized in RAID 0 arrays, each with their own individual device file, such as /dev/md0 and /dev/md1. These devices are then mirrored. This combines the speed advantages of RAID 0 with the data redundancy associated with mirroring. There are variations: for example, RAID 01 stripes two sets of RAID 1 mirrors. RAID 50 provides a similar combination of RAID 0 and RAID 5.

On the Job

Hardware RAID systems should be hotswappable. In other words, if one disk fails, the administrator can replace the failed disk while the server is still running. The system will then automatically rebuild the data onto the new disk. Since you can configure different partitions from the same physical disk for a software RAID system, the resulting configuration can easily fail if you use two or more partitions on the same physical disk. Alternatively, you may be able to set up spare disks on your servers; RAID may automatically rebuild data from a lost hard drive on properly configured spare disks.

RAID in Practice

RAID is associated with a substantial amount of data on a server. It's not uncommon to have a couple dozen hard disks working together in a RAID array. That much data can be rather valuable.

Inside the Exam

Creating RAID Arrays

During the Installation and Configuration portion of the Red Hat exams, it's generally easier to do as much as possible during the installation process. If you're asked to create a RAID array, it's easiest to do so with Disk Druid, which

works only during installation. You can create RAID arrays once RHEL is installed, but as you'll see in the following instructions, it is more time consuming and involves a process that is more difficult to remember.

However, if you're required to create a RAID array during your exam and forget to create it during the installation process, not all is lost. You can still use the tools described in this chapter to create and configure RAID arrays during the exam. And the skills you learn here can serve you well throughout your career.

If continued performance through a hardware failure is important, you can assign additional disks for failover, which sets up spare disks for the RAID array. When one disk fails, it is marked as bad. The data is almost immediately reconstructed on the first spare disk, resulting in little or no downtime.

Reviewing an Existing RAID Array

If you created a RAID array during the installation process, you'll see it in the /proc/ mdstat file. For example, I see the following on my system:

```
# cat /proc/mdstat
```

Certification Objective 8.06-Advanced Partitioning: Logical Volume Management

Logical Volume Management (LVM) (also known as the Logical Volume Manager) can allow you to manage active partitions. Before LVM, you had no easy way to increase or reduce the size of a partition after Linux was installed. With LVM2, you can even create read-write snapshots; but this is not part of the current exam requirements, so this book won't be addressing that feature.

For example, if you find that you have extra space on the /home directory partition and need more space on your /var directory partition for log files, LVM will let you reallocate the space. Alternatively, if you are managing a server on a growing network, new users will be common. You may reach the point at which you need more room on your /home directory partition. With LVM, you can add a new physical disk and allocate its storage capacity to an existing /home directory partition.

On the Job

While LVM can be an important tool to manage partitions, it does not by itself provide redundancy. Do not use it as a substitute for RAID. However, you can use LVM in concert with a properly configured RAID array.

Whenever you change an active PV, LV, or VG, unmount the volume first. If it's an essential filesystem such as the top-level root (/) directory, you may need to use **linux rescue** mode or a third-party bootable Linux such as Knoppix.

In essence, to create a new LVM system, you need to create a new PV, using a command such as pvcreate, assign the space to a VG with a command such as vgcreate, and allocate the space from some part of available VGs to an LV with a command such as lvcreate.

Inside the Exam

Logical Volume Management

One of the critical decisions during the Installation part of the RHCE and RHCT exams is whether you install in text or graphical mode. Text mode is faster. However, if you're required to create an LVM group during your exam, you can configure custom LVM groups with Disk Druid *only* it if you install RHEL in graphical mode.

I can't give you a concrete time savings between graphical and text mode; it depends on the traffic demands (how many other users) and the hardware available during your exam. I can say that when I installed the standard RHEL server configuration in graphical mode, it took 5 minutes longer than the same process in text mode. If your computer has more than 256MB of RAM (and more than 16MB of video memory), I suspect the difference would decrease.

If you forget to configure LVM during installation or are required to make changes, you can use the techniques I describe in this section to configure LVM groups after installation. Remember that the Red Hat Exam Prep guide suggests that RHCEs, during the Troubleshooting and System Maintenance portion of their exams, need to know how to

- Add, remove, and resize logical volumes.

To add space to an existing LVM system, you need to add free space from an existing VG with a command such as

lvextend. If you don't have any existing VG space, you'll need to add to it with unassigned PV space with a command such as vgextend. If all of your PVs are taken, you may need to create a new PV from an unassigned partition or hard drive with the pvcreate command.

You can also do much of this with the associated Logical Volume Management tool described near the end of this chapter.

Creating a Physical Volume

The first step in creating an LVM is to start with a physical disk. If you have a freshly installed hard disk, you can set up a PV on the entire disk. For example, if that hard disk is attached as the third PATA hard disk (/dev/hdc), and you haven't configured partitions on the drive, you'd run the following command:

```
#
# pvcreate /dev/hdc
```

Certification Summary

Kernels are at the heart of every operating system. The Linux kernel can be customized in a wide variety of ways. Two methods are based on loadable modules and changes to runtime parameters in the /proc directory.

The easiest way to update a kernel is to install (and not upgrade) from a Red Hat RPM, or use the yum command to install from your assigned repository. When you do, it automatically updates your boot loader files as needed.

Alternatively, the kernel can be optimized for your particular installation and hardware, and you have detailed control over its configuration. Once customized, it's a lot easier than it used to be to make a modular kernel; the **make rpm** command is all you need to create a customized kernel RPM. It's not perfect, as you still need to create your own initial RAM disk and corresponding stanza in your GRUB configuration file.

While it's best to configure RAID and LVM during the installation process, you may also need to do so after installation. RHEL supports a variety of software RAID types, including RAID 0, RAID 1, RAID 5, and RAID 6. LVM makes it easier to expand the size allocated to an existing filesystem; RHEL includes the GUI LVM tool to help.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 8.

The Basics of the Kernel

? The kernel lies at the heart of the operating system.

? Modular kernels, with separate device modules, are more efficient than monolithic kernels, where device drivers are integrated into the kernel.

? If you're going to update your kernel, you should keep a copy of your current working kernel.

? Linux kernels and related files are stored in the /boot directory.

? The /proc directory provides a window to what Linux sees in your computer.

? The **lsmod** command lists currently loaded modules; important related commands are **insmod**, **rmmod**, and modprobe.

? Basic RHEL modules can be loaded from the kernel-devel RPM.

New Kernels, the Easy Way

? Kernel version numbers are organized in *major.minor.patch* format. Red Hat adds a build number to the Linux kernels that it builds from source code.

? It's fairly easy to install a Red Hat kernel from RPM, as long as you remember to install and not upgrade. This allows you to return to the current working kernel if you have a problem.

? Sometimes all you need is a kernel patch, which supports upgrades of one patch version number. Unfortunately, patches are not always the best option for Red Hat built kernels.

? When you install a Red Hat kernel from RPM, the process should automatically update your GRUB boot loader.

Kernel Sources

? Kernel sources can be loaded from the kernel source RPM or from a Linux kernel tarball downloaded from a site such as ftp.kernel.org.

? Installing the Red Hat kernel source RPM requires the rpmbuild command.

? Once installed, the kernel source tree is available through /usr/src/redhat/BUILD directory.

Recompiling a Kernel

? To optimize the Linux kernel, it is a best practice to compile kernels with only needed elements and configure modules for most hardware.

? Your current kernel configuration is stored in the config-`uname -r` file in the /boot directory.

? You can modify kernel settings from the kernel source code directory with tools that you can open with one of the following commands: **make config**, **make menuconfig**, **make xconfig**, or **make gconfig**.

? Once you've made the proper backups and boot disks and set the **EXTRAVERSION** variable in your Makefile, you're ready to customize your kernel.

? Once you've settled on and saved your changes, run the **make rpm** command. It should compile your new kernel and create an RPM in the /usr/src/redhat/RPMS directory, which you can use to install your custom kernel.

Advanced Partitioning: Software RAID

? Red Hat supports several levels of software RAID, including RAID 0, RAID 1, RAID 5, and RAID 6.

? To make software RAID work, you need to designate the partition specifically as such in fdisk or **parted**.

? RAID arrays as configured are shown in /proc/mdstat.

? RAID arrays can be created and modified with the mdadm command.

Advanced Partitioning: Logical Volume Management

? LVM is based on physical volumes, logical volumes, and volume groups.

? You can create and add LVM systems with a wide variety of commands starting with **pv***, **lv***, and **vg***.

?

The GUI LVM tool is an alternative for those who don't remember all of the commands required to manage logical volumes.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions

The Basics of the Kernel

1. You are troubleshooting someone else's computer and are not sure what network card is inside it. You have checked the output from the dmesg command, but no network cards are listed, and even though you have a bunch of compiled network modules, none are currently loaded. What command might load the unknown network device? ?

2. What directory includes dynamic kernel configuration settings? ?

Answers

1. Any of these commands might load the unknown network device: **modprobe -lt net**, modprobe, and **modprobe eth0**.

2. The directory that includes dynamic kernel configuration settings is /proc.

New Kernels, the Easy Way

3. When you install a newer kernel from a Red Hat RPM, what else do you need to do before rebooting? ?

4. What happens if you use the **-U** switch to install a later Red Hat kernel RPM? ?

Answers

3. When you install an updated kernel from a Red Hat RPM, you shouldn't have to do anything else. As configured, it should automatically create an initial RAM disk in the /boot directory and add an appropriate stanza to the GRUB configuration file, /boot/grub/grub.conf.

4. If you use the **rpm -U** command, instead of **rpm -i**, you overwrite the currently active kernel. Since you're running from that kernel, you know that the active kernel works. If you overwrite it with a new kernel, and that new kernel doesn't work, you're out of luck.

Kernel Sources

5. Your kernel is version 2.6.19, and your architecture is x86_64. When you install and build the source code package associated with the Red Hat kernel, you'll find it in what directory?

____

6. When you install the kernel source code, you'll find kernel-2.6.spec in the /usr/src/redhat/SPECS directory. What command would you run to unpack the source code?

____

Answers

5. If your kernel is some version 2.6.19, and your architecture is x86_64, the kernel source code when installed can be found in the /usr/src/redhat/BUILD/kernel-2.6.19/linux-2.6.19.x86_64 directory.

6. When you install the kernel source code, with the kernel-2.6.spec in the /usr/src/redhat/SPECS directory, you can navigate to that directory and then unpack the code with the **rpmbuild -bb kernel-2.6.spec** command.

Recompiling a Kernel

7. Once you've navigated to the directory with the kernel source code, name two commands that would open a menu that can help you configure a custom kernel. Assume all necessary dependencies are also installed.

____

____

8. You have just compiled and installed a custom kernel from kernel.org. What two other things do you need to do before you restart and boot from that kernel?

____

____

Answers

7. Once you've navigated to the directory with the kernel source code, there are not just two, but four commands available that open a menu that can help you configure a custom kernel: **make config**, **make menuconfig**, **make xconfig**, and **make gconfig**.

8. You have just compiled and installed a custom kernel RPM. The two other things do you need to do before you restart and boot from that kernel is create an initial RAM disk with the **mkinitrd** command and create a new stanza in the GRUB configuration file, /boot/grub/grub.conf.

Advanced Partitioning: Software RAID

9. In what file can you find the current software RAID configuration?

____

10. What command can help you build a new RAID array? Just list the basic command; assume that you'll go to the man page for more information.

Answers

9. You can find the current software RAID configuration in /proc/mdstat.

10. The mdadm command can help you build a new RAID array.

Advanced Partitioning: Logical Volume Management

11. Once you've created a new partition and set it to the Logical Volume Management filetype, what command adds it as a PV?

---

12. Once you've added more space to an LV, what command would expand the formatted filesystem to fill the new space?

Answers

11. Once you've created a new partition and set it to the Logical Volume Management filetype, the command that adds it as a PV is pvcreate. For example, if the new partition is /dev/hdb2, the command is **pvcreate /dev/hdb2**.

12. Once you've added more space to an LV, the command that would expand the formatted filesystem to fill the new space is resize2fs.

Lab Questions

The Red Hat exams are unique based on their reliance on labs and hands-on demonstrations. With these questions, you're practicing the skills you need on both Red Hat exams.

Lab 1

1. In this lab, you'll install an updated kernel. This requires that an updated kernel be available, and you may not be able to run this lab unless you have such a kernel. However, if an older kernel is available for the distribution, you could add the **--force** switch to the **rpm** command used in this lab to downgrade to the old kernel temporarily.

In most cases, if you have an appropriate connection to the RHN or a rebuild repository (or even the update repositories for Fedora Core 6), you should be able to install the new kernel. List two methods you could use.

Finally, observe what happens to /boot/grub/grub.conf. What changed? What's the default kernel?

Answers

1. Installing an updated kernel is easy. Assuming it's available from a repository such as the RHN, all you should have to do is run the following command:

```
#
# yum install kernel
```

Lab 2

2. In this lab, you'll customize and recompile the kernel. Yes, this is beyond what is listed in the Red Hat Exam Prep guide. However, recent changes to the kernel compilation process leads me to believe that Red Hat may reintroduce this skill into the RHCE exam sometime in the near future.

This lab is more of a detailed kernel building exercise than a typical lab. Even if you don't need to compile a kernel on the Red Hat exams, I predict that you will need to compile a kernel at some point as a Linux system administrator. See the Lab Answers section at the end of this chapter for the exercise.

If you want to try things out for yourself, download and install the kernel src.rpm for the active kernel. Install it, and run the rpmbuild command on the associated .spec file. Customize some minor part of the kernel, change the **EXTRAVERSION** directive in the associated Makefile, and run **make rpm** to compile and create an RPM for the custom kernel.

Answers

2. Before building a new kernel, you need the correct RPM packages (with dependencies), often associated with the Development Tools package group. As with the rest of this chapter and the Red Hat exams, this assumes that you have a PC with a 32-bit Intel type CPU. The procedures for other CPUs vary and are not, as of this writing, covered on the Red Hat exams.

The following list of RPMs are associated with the source code (version numbers omitted). This does not include dependencies. When possible, use the yum command to satisfy and install these dependencies automatically.

```
k
ke rne l- u name - r`s rc .rpm
u
u ni de f
 r
 rpm- bui ld
g
g li b- hea de rs
g
g li b- de ve l
c
```

Lab 3

3. In this lab, you'll add a new RAID array to your system. This is possible if you have free space on the current drive or can add a new drive. That's not as difficult as it appears if you have a VMware-based system. You can even use a USB key for this purpose (assuming you don't need to save any data on that key).

One way to test this lab is to create a separate RAID partition for the /tmp directory. Unless you use it for downloads, it's fairly small; on my system, it has 120KB of data. (If you're using /tmp for downloads, chances are you're using a production system, and that can be dangerous, unless your backups are in order.)

Create a simple RAID 1 mirror. All you need are two partitions. While it's best if the partitions are on different physical drives, it's not required for this lab. Remember to activate the array. Assign and mount it on a new directory. Copy files from your home directory. If the space used by your home directory is greater than available from the RAID array, you don't have to copy everything. Add the information to your /etc/fstab, and test the result.

Answers

3. This lab assumes that you have some unpartitioned space, either on an existing drive or a new drive that you've just added. A VMware machine is one excellent option that makes adding a new drive relatively easy. For the purpose of this lab, even a USB key will suffice.

To add a new partition, you'll need to use either fdisk or **parted**. Allocate no more than half of the available free space to the first partition. Make sure to designate the RAID filetype. If you use fdisk, you'll need to either reboot or run partprobe to make Linux reread the new partition table. Whatever you do, one simple way to make sure you have a new partition of the RAID filetype is with the following command:

```
#
#  fdisk - l
```

Lab 4

4. Lab 4 and Lab 5 are to be run in sequence. ?

In Lab 4, you'll add a new LV to your system. It doesn't have to be complex; you can use free extents from an available LVM or a single new LVM partition. If possible, you'll want to leave some free extents for Lab 5, or you'll have to create another new LVM partition in that lab.

One way to test this lab is to create a separate LVM for the /tmp directory. Unless you use it for downloads, it's fairly small; on my system, it has 120KB of data. (If you're using /tmp for downloads, chances are you're using a production system, and that can be dangerous, unless your backups are in order.)

Add the information to your /etc/fstab, and run **mount -a** to test the result.

Answers

4. For this lab, I assume that you've never created an LV before on the local system. The first step is to create (or reallocate) dedicated partitions for this purpose and assign it to the LVM file type. In fdisk, that's file type **8e**; to set the first partition to LVM in **parted**, run **set 1 lvm on**. Repeat the process for additional LVM partitions. For the purpose of this lab, I've created a single partition, /dev/sdb1, with 200MB, and assume you're working from the command line. Of course, you can do everything in this lab-after creating a partition-with the GUI Logical Volume Management tool.

Next, create a PV; for example, if the new LVM partition you created is /dev/sdb1, run the following command:

```
#
# pvcreate /dev/sdb1
```

Lab 5

5. In Lab 5, you'll increase the size of the new LV. First, run the **df** command to review the space taken by the current version of the LV. Unmount it from the directory you've configured. Use available free space from existing LVM partitions, or create a new LVM partition. Remount the newly expanded LV, and test the result. Are the files still there? What is the size of the new LV? ?

Answers

5. In this lab, you'll use free VG space from Lab 4. If you don't have that available, you'll have to repeat the steps described in Lab 4 to create a new PV and VG. To confirm available VG space, run the **vgs** command. In my case, I see the following output from my system:

```
#
# vgs
 V
 VG          #PV #LV #SN Attr    VSize    VFree
```

Chapter 9: Apache and Squid

Overview

Unix was developed by AT&T in the late 1960s and early 1970s, and it was freely distributed among a number of major universities during those years. When AT&T started charging for Unix, a number of developers tried to create clones of this operating system. One of these clones, Linux, was developed and released in the early 1990s.

Many of these same universities were also developing the network that evolved into the Internet. With current refinements, this makes Linux perhaps the most Internet-friendly network operating system available. The extensive network services available with Linux are not only the tops in their field, but they create one of the most powerful and useful Internet-ready platforms available today at any price.

Currently, Apache is the most popular Web server on the Internet. According to the Netcraft (www.netcraft.com) survey, which tracks the Web servers associated with virtually every site on the Internet, Apache is currently used by more Internet Web sites than all other Web servers combined. Apache is included with RHEL 5.

RHEL 5 also includes a number of other network services. The service that also focuses on Web access is the Squid Proxy Server, which caches frequently used pages locally.

This chapter deals with the basic concepts surrounding the use of these services and a basic level of configuration. In all cases, the assumption is that your network settings are correct and functioning properly. If you're having problems with your network configuration, read Chapter 7.

As for the RHCE exam, you may have to configure or troubleshoot either Apache or Squid. So as you read this chapter and look through the configuration files and exercises, be willing to experiment. And practice, practice, practice what you learn.

Certification Objective 9.01-The Apache Web Server

Apache is by far the most popular Web server in use today. Based on the HTTP daemon (**httpd**), Apache provides simple and secure access to all types of content using the regular HTTP protocol as well as its secure cousin, HTTPS.

Apache was developed from the server code created by the National Center for Supercomputing Applications (NCSA). It included so many patches that it became known as "a patchy" server. The Apache Web server continues to advance the art of the Web and provides one of the most stable, secure, robust, and reliable Web servers available. This server is under constant development by the Apache Software Foundation (www.apache.org).

Inside the Exam

This chapter directly addresses two items in the Red Hat Exam Prep guide. This is the first chapter to focus on network services, as required of RHCE candidates. Per the latest Exam Prep guide, RHCT candidates do not need to be too concerned with this chapter. As noted in the Exam Prep guide, RHCE candidates "must be capable of configuring the following network services" during the Installation and Configuration portion of that exam:

- HTTP/HTTPS
- Web Proxy

Although you can use a number of different packages to configure HTTP, HTTPS, and Web Proxy services, the publicly available RH300 course outline focuses these services on Apache as a regular and secure Web server, and Squid as the Web Proxy server. The Exam Prep guide also notes that RHCEs should be able to

- Diagnose and correct problems with network services.
- Diagnose and correct networking services problems where SELinux contexts are interfering with proper operation.

This includes those services listed in the Installation and Configuration portion of the RHCE exam. For every network service, you also need to

- Install the packages needed to provide the service.
- Configure SELinux to support the service.
- Configure the service to start when the system is booted.
- Configure the service for basic operation.
-

Configure host-based and user-based security for the service.

Installing the required packages is trivial. You'll make sure the service is started when the system is booted with the appropriate chkconfig commands. Most of this chapter is dedicated to configuring the service for *basic* operation. Some services support host-based and user-based security in their configuration files; others support it with the tools described in Chapter 15. SELinux is also most easily configured using the SELinux Management tool described in Chapter 15.

While there are numerous other Web servers available, Apache is the only Web service described in the current RH300 course outline.

Apache is a service; basic Apache clients are Web browsers. Therefore, only those concerned with the RHCE need to read this chapter. This provides the briefest of overviews on Apache. For more information, read the documentation online at http://httpd.apache.org/docs-2.2.

Apache 2.2

Red Hat Enterprise Linux includes the latest major release of Apache, which is 2.2.x as of this writing. While there are major differences from previous versions of Apache (1.3.x, 2.0.x), if you're a Web administrator or developer, the differences with respect to the RHCE exam are fairly straightforward. The current version supports virtual hosts and access control, as well as secure (HTTPS) Web services. If you're interested in more, a full list of new features is available from http://httpd.apache.org/docs/2.2/new_features_2_2.html.

The following cites a few of the major changes:

- **New packages** If you're installing Apache from the Red Hat Installation RPMs, all the package names have changed. As you'll see in the following section, most start with *httpd*. Strangely enough, the username associated with Apache services is now *apache*.
- **Modular directive files** Basic directives, such as those based on Perl, PHP, or the Secure Socket Layer, are now configured separately in the /etc/httpd/conf.d directory. They are automatically included in the Apache configuration with the following directive in /etc/httpd/conf/httpd.conf:

Certification Objective 9.02-Apache Access Configuration

There are several parameters associated with security on the Apache Web server. The security of the server is enforced in part by firewalls and SELinux. Internal Apache security measures are associated with the main Apache httpd.conf configuration file.

Now that you've glanced at the configuration file, it's time to analyze it, and its associated directories, with a view toward security.

Basic Apache Security

You can modify the httpd.conf configuration file to secure the entire server or manage security on a directory-by-directory basis. Directory controls secure access by the server, as well as users who connect to the Web sites on the server. To explore the basics of Apache security, start with the first default active line in httpd.conf:

```
S
ServerTokens OS
```

Certification Objective 9.03-Virtual Hosts

Another useful feature of Apache 2.2 is its ability to manage Web sites using a single IP address. You can do so by creating multiple virtual hosts on the same Web server. You can configure virtual hosts for regular Web sites in the main Apache configuration file, /etc/httpd/conf/httpd.conf. In that way, you can link multiple domain names such as www.example.com and www.mommabears.com to the same IP address on the same Apache server.

On the Job

The example.com, example.org, and example.net domain names cannot be registered and are officially reserved by the Internet society for documentation.

You can also create multiple secure Web sites that conform to the HTTPS protocol by configuring virtual hosts in the /etc/httpd/conf.d/ssl.conf configuration file. While the details vary, the basic directives that you'd use in this file are the same.

Exam Watch

While truly secure HTTPS sites include server certificates, there is no cited requirement in the Red Hat Exam Prep guide or associated RH300 course to create such certificates.

Virtual Hosts

As described earlier, Section 3 of the default httpd.conf includes sample commands that you might use to create one or more virtual hosts. To activate the virtual host feature, you'll first want to activate this directive:

```
#
# NameVirtualHost *:80
```

Certification Objective 9.04-The Squid Web Proxy Cache

Squid is a high-performance HTTP and FTP caching proxy server. It is also known as a Web proxy cache. It can make your network connections more efficient. As it stores data from frequently used Web pages and files, it can often give your users the data they need without their systems having to look to the Internet.

Studies on very busy networks suggest that a Squid server can reduce the size, or bandwidth, of your Internet connection by 10 to 20 percent. That can lead to considerable savings for larger offices.

Squid uses the Inter-Cache Protocol (ICP) for transfers between participating peer and parent/child cache servers. It can be used either as a traditional caching proxy or as a front-end accelerator for a traditional Web server. Squid accepts only HTTP requests but speaks FTP on the server side when FTP objects are requested. For more information, see www.squid-cache.org. One book dedicated to this service is Duane Wessels's *Squid: The Definitive Guide*, published by O'Reilly.

Key Squid Files and Directories

The Squid RPM package is installed by default when you install the Web Server package group. So if you've installed Apache and have not tinkered with the defaults, the Squid RPM should also be installed on your computer. This RPM package installs a substantial number of files and scripts; some of the key files include the following:

- **/etc/rc.d/init.d/squid** Start/stop script
- **/etc/squid/** Configuration directory
- **/etc/sysconfig/squid** Other configurable options
- **/usr/share/doc/squid-*versionnumber*** Documentation, mostly in HTML format
- **/usr/lib/squid/** Support files and internationalized error messages
- **/usr/sbin/squid** Main Squid daemon
- **/usr/share/squid** Various squid configuration add-ons
- **/var/log/squid/** Log directory
- **/var/spool/squid/** Cache directory (once Squid is active, this directory includes hundreds of MBs, and maybe more, in hashed directories)

Starting Squid on Reboot

The Squid Web Proxy is not started by default. To do so, you'll want to activate it using a command such as chkconfig or the Service Configuration utility described in Chapter 3. The easiest way to set Squid to start the next time you boot Linux is with the following command:

```
#
# chkconfig squid on
```

Certification Summary

You can configure a number of network sharing services on your RHEL computer. Apache is the most important Web server on the Internet. Squid allows you to save bandwidth.

Apache was developed from the NCSA Web server. Once the appropriate packages are installed, you can access a structure and sample Web pages in the /var/ www/html directory, based on the /etc/httpd/conf/httpd.conf configuration file. The httpd.conf file is organized in containers. You can create virtual hosts for multiple Web sites on your computer, even if you have only one IP address.

Squid is a proxy server that allows a network to filter its HTTP and FTP traffic through a cache. Requests are taken from the cache when possible. This reduces the load between the LAN and the Internet, reducing your network costs. When users access cached files, they get better performance from the external network.

As the RHCE is a performance-based exam, it is important to practice all the skills discussed in this chapter. You may need to use these skills on the exam!

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 9.

The Apache Web Server

? Red Hat Enterprise Linux includes the Apache Web server, which is currently used by more than twice as many Internet Web sites as all other Web servers combined.

? Web services are an easy way to provide simple, secure access to many types of documents. The Apache Web server provides both normal and secure Web services using the HTTP and HTTPS protocols.

Apache Access Configuration

? Apache supports security by specifying active ports through the **Listen** and **NameVirtualHost** directives.

? Apache supports host-based security by IP address or domain name.

? Apache supports user-based security by password, with the help of the htpasswd command.

Virtual Hosts

? With Apache 2.2, you can configure multiple Web sites on your server, even if you have only one IP address. This is possible through the use of virtual hosts.

? The RHEL configuration supports the configuration of virtual hosts for regular Web sites at the end of the /etc/httpd/conf/httpd.conf file.

? The RHEL configuration supports the configuration of secure virtual hosts for regular Web sites at the end of the /etc/httpd/conf.d/ssl.conf file.

The Squid Web Proxy Cache

? Squid is a high-performance HTTP and FTP caching proxy server.

? The main Squid configuration file is long, but all you need to do in /etc/ squid/squid.conf is configure the following parameters: **visible_hostname**, **http_access**, and **acl**.

?

Squid can refer requests to sibling and parent proxy servers. If the request still isn't available, a parent proxy server refers the request to the Internet.

?

Once Squid is configured, you can set each computer on the LAN to browse Web pages to the proxy server on port 3128, or redirect requests with the help of an appropriate iptables routing command.

PREV NEXT

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

The Apache Web Server

1. What is the Apache directive that specifies the base directory for configuration and log files?

   ____

2. Once you've modified httpd.conf, what command should you use to reread this file?

   ____

3. What directive specifies the TCP/IP port associated with Apache?

   ____

Answers

1. The **ServerRoot** directive sets the default directory for the Apache server. Any files and directories not otherwise configured-or configured as a relative directory-are set relative to **ServerRoot**.

2. There are two basic ways to make Apache reread the configuration file. You can restart the service with a command such as **apachectl restart**, or you can keep Apache running and make it reread the file with **service httpd reload**.

3. The **Listen** directive specifies the TCP port associated with Apache.

Apache Access Configuration

4. What command creates the /etc/httpd/passwords file and configures a password for user elizabeth?

   ____

5. If you see the following directives limiting access within the stanza for a virtual host, what computers are allowed access?

   ```
   O
   Order Allow,Deny
   A
   ```

6. What standard ports do you need to open in a firewall to allow access to a regular and secure Web site? ____

Answers

4. The command that creates the /etc/httpd/passwords file and configures a password for user elizabeth is **htpasswd -c /etc/httpd/passwords elizabeth**. If /etc/httpd/passwords already exists, all that's required is **htpasswd elizabeth**.

5. As described in the chapter, the **Order Allow,Deny** directive denies access to all systems by default, except those explicitly allowed access. So access is limited to computers on the 192.168.0.0/24 network.

6. The standard ports you need to open in a firewall to allow access to a regular and secure Web sites are 80 and 443.

Virtual Hosts

7. What file does RHEL provide to help you configure a virtual host as a secure server? ____ ?

8. If you're creating an IP address-based virtual host, how many IP addresses would you need for three virtual servers? _______________ ?

9. If you want to check your configuration of virtual hosts, what switch can you use with the **httpd** command? _______________________________ ?

Answers

7. The file associated with secure servers for virtual hosts is ssl.conf in the /etc/httpd/conf.d/ directory.

8. Three IP addresses are needed for three virtual servers. IP address-based virtual hosts each require their own IP addresses. If you want to configure multiple virtual hosts, you need name-based virtual hosts, typically starting with the **NameVirtualHost *:80** directive.

9. To check your configuration of virtual hosts, you can use two switches with the **httpd** command: **httpd -S** checks the configuration file, including virtual host settings. Alternatively, **httpd -D DUMP_VHOSTS** focuses on the virtual host configuration.

The Squid Web Proxy Cache

10. Which directive in squid.conf is used to point to the network IP address served by Squid? ____ ?

11. What command configures the Squid cache directories and is required before you start this service for the first time? ___________________________________ ?

12. What port and protocol are the defaults for Squid communication? _____________ ?

Answers

10. The **acl** and **http_access allow** directives in squid.conf can be used to specify the network to be served by Squid. For example, the following directives support access from the 192.168.0.0/24 network:

```
a
acl local_net src 192.168.0.0/24
 h
```

11. The **squid -z** command sets up the directories for the Squid cache, which is required before you start the Squid service for the first time.

12. The default port for Squid communication is 3128; the default protocol is ICP (InterCache Protocol).

PREV NEXT

Lab Questions

Lab 1

1. In this first lab, you'll install and configure Apache to start and run automatically the next time you boot your computer. You'll also configure the default home page for the local distribution as the default home page for the local computer. ?

Answers

1. First, make sure the Apache Web server is installed. If an **rpm -q httpd** command tells you that it is missing, you haven't installed the Web Server package group. The most efficient way to do so is with the **yum install "Web Server"** command. (To find appropriate package group names, see the comps-rhel5-server-core.xml file in the Server/repodata subdirectory on the first installation CD, or run the **yum groupinfo** command.) This assumes a proper connection to a repository; if you're using RHEL 5, this requires a properly enabled subscription to the Red Hat Network. Other update options are covered in Chapter 5.

   To configure Apache to start, run the **apachectl start** command. To make sure it starts the next time you boot your computer, run the **chkconfig httpd on** command.

   Once Apache is installed, you should be able to access it by opening a browser and navigating to http://localhost. You can see in the default Apache configuration file that the **DocumentRoot** is located in /var/www/html. The default home page is located at /usr/share/doc/HTML/index.html. You can copy that index.html file to the /var/www/html directory and test the result by navigating once again to http://localhost. If you did not copy the other files associated with the default home page, you'll be missing some icons.

Lab 2

2. In this lab, you'll configure two Web sites on the local Apache server. Call them big.example.big and small.example.small. Don't forget to create the directories that you need and set up these Web sites on your DNS server or /etc/hosts file. Make sure your Web sites are accessible to users from remote computers on your network. Add an appropriate index.html file to the **DocumentRoot** for each Web site. Simple Web pages, such as a single line of text, are acceptable (no HTML coding is necessary). Don't forget that SELinux settings need to be compatible with what you configure. ?

Answers

2. This lab requires that you create two virtual hosts in the main Apache configuration file, /etc/httpd/conf/httpd.conf. To make this happen, you should take the following steps:

1. The **ServerRoot** directive sets the default directory for the Apache server. Any files and directories not otherwise configured-or configured as a relative directory-are set relative to **ServerRoot**. Don't change this unless you're ready to adjust the SELinux contexts of the new directory accordingly.
2. Set the **NameVirtualHost** directive to the port (80) serving your intended network audience. Don't assign any IP addresses.
3. Add a **VirtualHost** container with the same IP address.
4. Assign the **ServerAdmin** to the e-mail address of this Web site's administrator.
5. Configure a unique **DocumentRoot** directory.
6. Set the first **ServerName** to big.example.big.
7. Add **ErrorLog** and **CustomLog** directives, and set them to unique file names in the /etc/httpd/ logs directory. With the default **ServerRoot**, you can use a relative logs directory, such as this:

Lab 3

3. Continuing on with Apache, now configure secure versions for each of your two Web sites. Make sure that appropriate directories are available for each secure Web site.

Answers

3. The basics of this lab are straightforward. You'll need to repeat the same basic steps that you performed in Lab 2; you're just editing the /etc/httpd/conf.d/ssl.conf configuration file. However, you should be concerned about the following:

    1. Make sure that the top **VirtualHost** directive points to the IP address that you're using for your Web server.
    2. Set up the **DocumentRoot** in a directory different from a regular Web server.
    3. Configure the **ErrorLog** and **CustomLog** separately; select names to associate these log files with the name of the secure Web site.
    4. Continuing on with Apache, now configure secure versions for each of your two Web sites. Make sure that appropriate directories are available for each secure Web site.

Lab 4

4. Set up a Squid Proxy Server on your computer. Set up access to your LAN on the 10.11.12.0/255.255.255.0 network. Assign appropriate values to **acl**, **http_access**, and **visible_hostname**. Set up the cache directories for Squid. Make sure it starts now and automatically the next time you reboot your computer. If there are problems, make sure the SELinux settings are compatible.

Answers

4. Squid is installed by default when you install the Web Server package group. To configure a Squid Proxy Server for your network, you'll need to configure /etc/squid/squid.conf. Assume the name of your computer is myproxy, and you're arbitrarily assigning mylan as the name for your LAN. If your network IP address is not 10.11.12.0, substitute accordingly. In this file, you'll need to add directives similar to:

```
visible_hostname=myproxy
acl mylan src 10.11.12.0/255.255.255.0
```

Chapter 10: Network File-Sharing Services

Overview

Linux is designed for networking, and there are three major protocols associated with sharing files on a network: NFS, FTP, and Samba. While some excellent GUI tools are available, I recommend that you learn to configure these services from the command line. If you know these services, you can do more in less time by directly editing key configuration files.

This chapter starts with a description of the Network File System (NFS), a powerful and versatile way of sharing filesystems between servers and workstations. NFS clients are installed with a default installation of RHEL 5 and support connections to NFS servers. RHEL includes an excellent GUI-based configuration tool.

The chapter continues with the Very Secure FTP (vsFTP) daemon, which provides both basic and secure FTP server services. With vsFTP, you can secure users, directories, subdirectories, and files with various levels of access control.

This chapter finishes with a detailed analysis of networking with the various Microsoft Windows operating systems. Microsoft networking is based on the Common Internet File System (CIFS), which was developed from the Server Message Block (SMB) protocol. Samba was developed as a freely available SMB server for all Unix-related operating systems, including Linux, and has been upgraded to support CIFS.

Samba interacts with CIFS so transparently that Microsoft clients cannot tell your Linux server from a genuine Windows NT/2000/XP/2003/Vista server, and with Samba on Linux there are no server, client, or client access licenses to purchase. If you can learn to edit the main Samba configuration file from the command line interface, you can configure Samba quickly. RHEL includes a GUI alternative-the Samba Server Configuration utility. There's even a Samba Web Administration Tool, which we won't discuss in this book.

As you learn about these network services, you're learning about the services that you might configure and/or troubleshoot on the Red Hat exams. Take the time you need to understand the configuration files associated with each of these services, and practice making them work on your Linux computer. In some cases, two computers running Linux will be useful to practice what you learn in this chapter.

Certification Objective 10.01-Configuring a Network File System (NFS) Server

The NFS is the standard for sharing files and printers on a directory with Linux and Unix computers. It was originally developed by Sun Microsystems in the mid- 1980s. Linux has supported NFS (both as a client and a server) for years, and NFS continues to be popular in organizations with Unix- or Linux-based networks.

Inside the Exam

More Network Services

The Red Hat Exam Prep guide suggests that you can expect to configure NFS, FTP, and Samba servers during the Installation and Configuration portion of the RHCE exam. As described in the Red Hat Exam Prep guide, for each of these services, RHCEs must be able to do the following:

- Install the packages needed to provide the service.
- Configure SELinux to support the service.
- Configure the service to start when the system is booted.
- Configure the service for basic operation.
- Configure host-based and user-based security for the service.

It also suggests that you need to know how to "diagnose and correct problems with (these) network services" as well as SELinux-related network issues during the Troubleshooting and System Maintenance portion of this exam.

In this chapter, the NFS service should be installed automatically with RHEL 5. vsFTP and Samba are important server options for many administrators, and it's very possible that you'll install them during the RHCE exam. For each service, remember to use a command such as chkconfig to make sure it starts the next time you boot Linux. It also will help you get full credit for the work you do on the Red Hat exams.

You can create shared NFS directories directly by editing the /etc/exports configuration file, or with Red Hat's NFS Configuration tool. As NFS servers come first, that's what this chapter covers first.

Exam Watch

I believe even beginning Linux administrators should know how to connect to a shared NFS directory. While the Red Hat Exam Prep guide does not explicitly require RHCT candidates to have this knowledge, it is consistent with the spirit of that exam. The Red Hat Exam Prep guide does explicitly require that RHCE candidates know how to configure and troubleshoot an NFS server.

NFS Server Configuration and Operation

NFS servers are relatively easy to configure. All you need to do is export a filesystem, either generally or to a specific host, and then mount that filesystem from a remote client. In Chapter 2, you configured an NFS server to install RHEL over a network. This chapter goes further into the basics of NFS server configuration and operation.

Required Packages

Two RPM packages are closely associated with NFS: portmap and nfs-utils. They should be installed by default. Just in case, you can use the **rpm -q** ***packagename*** command to make sure these packages are installed. The **rpm -ql** ***packagename*** command provides a list of files installed from that package. The nfs-utils package includes a number of key files. The following is not a complete list:

- **/etc/rc.d/init.d/nfs** Control script for NFS, hard linked to /etc/init.d/nfs
- **/etc/rc.d/init.d/nfslock** Control script for lockd and statd, which locks files currently in use, hard linked to /etc/init.d/nfslock
- **/usr/share/doc/nfs-utils-*versionnumber*** Documentation, mostly in HTML format
- **Server daemons in /usr/sbin** rpc.mountd, rpc.nfsd, rpc.rquotad
- **Server daemons in /sbin** rpc.lockd, rpc.statd
- **Control programs in /usr/sbin** exportfs, nfsstat, nhfsgraph, nhfsnums, nhfsrun, nhfsstone, showmount
- **Status files in /var/lib/nfs** etab, rmtab, statd, state, xtab

The portmap RPM package includes the following key files (also not a complete list):

- **/etc/rc.d/init.d/portmap** Control script, hard linked to /etc/init.d/portmap
- **/usr/share/doc/portmap-4.0** Documentation
- **Server daemon in /sbin** portmap
- **Control programs in /usr/sbin** pmap_dump, pmap_set

Configuring NFS to Start

Once configured, NFS can be set up to start during the Linux boot process or can be started with the **service nfs start** command. NFS also depends on the portmap package, which helps secure NFS directories that are shared through /etc/exports. Because of this dependency, you need to make sure to start the portmap daemon before starting

NFS, and don't stop it until after stopping NFS.

Exam Watch

Remember that both the portmap and nfs daemons must be running before NFS can work.

The **nfs** service script starts the following processes:

- **rpc.mountd** Handles mount requests
- **nfsd** Starts an nfsd kernel process for each shared directory
- **rpc.rquotad** Reports disk quota statistics to clients

If any of these processes is not running, NFS won't work. Fortunately, it's easy to check for these processes by running the **rpcinfo -p** command. As with other service scripts, if you want it to start when RHEL boots, you'll need to run a command such as this:

```
#
# chkconfig nfs on
```

Certification Objective 10.02-Client-Side NFS

Now you can mount a shared NFS directory from a client computer. The commands and configuration files are similar to those used for any local filesystem. In the previous section, you configured an NFS server; for the moment, stay where you are, because computers that serve as NFS servers are also NFS clients.

Mounting an NFS Directory from the Command Line

Before doing anything elaborate, you should test the shared NFS directory from a Linux or Unix client computer. But first, you should check for the list of shared NFS directories. If you're on an NFS server and want to check the local list, the command is easy:

```
#
# showmount -e
```

Certification Objective 10.03-The File Transfer Protocol and vsFTPd

The File Transfer Protocol is one of the original network applications developed with the TCP/IP protocol suite. It follows the standard model for network services, as FTP requires a client and a server. The FTP client is installed by default on most operating systems, including Red Hat Enterprise Linux. If you've installed the FTP Server package group, you've installed the default Red Hat FTP Server, the very secure FTP (vsFTP) daemon.

In this section, you'll look solely at the vsFTP server. The lftp client was examined in Chapter 1, and other FTP servers are not supported on RHEL 5.

Installing the Very Secure FTP Server

The only FTP server included with RHEL is vsFTP. If it isn't already installed, you could use a GUI tool to install it. But the simplest method, based on a proper connection to the Red Hat Network or a rebuild repository, is with the following command:

```
#
# yum install vsftpd
```

Certification Objective 10.04-Samba Services

Microsoft's CIFS was built on the Server Message Block (SMB) protocol. SMB was developed in the 1980s by IBM, Microsoft, and Intel as a way to share files and printers over a network.

As Microsoft has developed SMB into CIFS, the Samba developers have upgraded Samba accordingly. Samba services provide a stable, reliable, fast, and highly compatible file and print sharing service that allows your computer to act as a client, a member server, or even a Primary Domain Controller (PDC) or a member of an Active Directory (AD) service on Microsoft-based networks. While Samba does not include every feature built into the latest Microsoft networks, I have confidence that it will in the near future.

On the Job

I look forward to the release of Samba 4.0, which will make it possible for Linux to act as an AD controller on a Microsoft-based network. However, I don't believe you will see Samba 4.0 in Red Hat distributions until RHEL 6 is released.

SMB network communication over a Microsoft-based network is also known as NetBIOS over TCP/IP. Through the collective works of Andrew Tridgell and the Samba team, Linux systems provide transparent and reliable SMB support over TCP/IP via a package known as Samba. You can do four basic things with Samba:

- Share a Linux directory tree with Windows and Linux/Unix computers
- Share a Windows directory with Linux/Unix computers
- Share a Linux printer with Windows and Linux/Unix computers
- Share a Windows printer with Linux/Unix computers

Samba emulates many of the advanced network features and functions associated with the Win9*x*/Me and NT/2000/XP/2003/Vista operating systems through the SMB protocol. Complete information can be found at the official Samba Web site at www.samba.org. It is easy to configure Samba to do a number of things on a Microsoft-based network. Here are some examples:

- Participate in a Microsoft Windows 9*x*-style workgroup or an NT/2000/XP/2003 domain as a client, member server, or even a PDC.Share user home directories.
- Act as a WINS (Windows Internet Name Service) client or server.
- Link to or manage a workgroup browse service.
- Act as a master browser.

- Provide user/password and share security databases locally, from another Samba server or from a Microsoft NT 4 PDC.
- Configure local directories as shared SMB filesystems.
- Synchronize passwords between Windows and Linux systems.
- Support Microsoft Access Control Lists.

Samba can do more, but you get the idea. Samba features are configured through one very big file, smb.conf, in the /etc/samba directory. As this file may intimidate some users, Red Hat's Samba Server Configuration tool (**system-config-samba**) provides an easier interface. RHEL 5 does not include the Samba Web Administration Tool, so don't expect it to be available on the Red Hat exams.

Exam Watch

I believe that Red Hat's Samba Server Configuration utility is an effective tool. But if you know how to edit the /etc/samba/smb.conf configuration file in a text editor, you're more likely to have time to configure the other elements you need to pass the exam. But don't be afraid to use the method that is fastest for you.

Installing Samba Services

If you selected the Windows File Server package group when you installed RHEL 5, the Samba RPM packages should already be installed. These are the four Samba RPM packages that you need:

- The samba RPM package includes the basic SMB server software for sharing files and printers.
- The samba-client RPM package provides the utilities needed to connect to shares from Microsoft computers.
- The system-config-samba package installs the Red Hat Samba Server Configuration utility.
- The samba-common RPM package contains common Samba configuration files.

It's easy to start the Samba Server Configuration tool. You can do so from a command line interface in the GUI with the **system-config-samba** command. Alternatively, you can choose System (or KDE Main Menu) | Administration | Server Settings | Samba. Either command opens the utility shown in Figure 10-3.

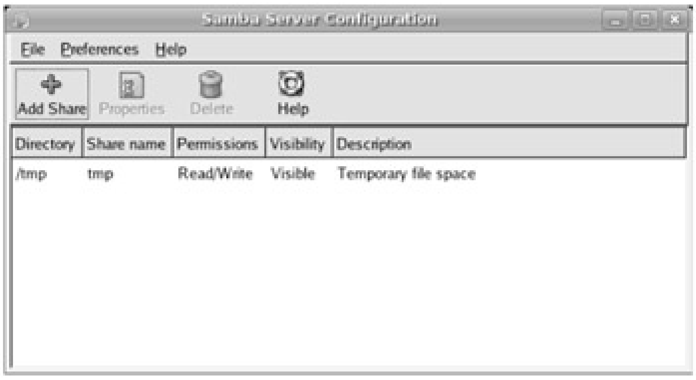

Figure 10-3: Samba Server Configuration utility

Configuring SELinux Support for Samba

There are four directives associated with making a Samba server work with SELinux in targeted mode:

- **allow_smb_anon_write** supports the writing of files to directories configured with the **public_content_rw_t** SELinux setting.
- **samba_enable_home_dirs** allows Samba to share users' home directories.
- **samba_share_nfs** allows Samba to share directories already shared via NFS.
- **use_samba_home_dirs** supports remote access to local home directories using Samba.

To set these directives, use the **setsebool** command. For example, to allow Samba to share local home directories with others on the network, run the following command:

```
#
# setsebool -P samba_enable_home_dirs 1
```

The Red Hat Samba Server Configuration Utility

RHEL includes Red Hat's graphical configuration tool for Samba, **system-config-samba**, which you can install from the RPM of the same name. Before you use this tool to modify your configuration, back up the files in your /etc/samba directory.

Also known as the Samba Server Configuration utility, you can use this tool to set basic global parameters and configure shared directories. You can start it from a GUI command line with the **system-config-samba** command, or you can choose System (or KDE Main Menu) | Administration | Server Settings | Samba. You saw the basic tool back in Figure 10-3.

This tool is straightforward. You can configure general Samba directives such as **security level** and **workgroup** through the Preferences | Server Settings command. The Add button enables you to set up a new share.

You can also use this tool to configure Samba usernames and passwords. In other words, you can use this tool to configure your smb.conf file as well as Samba usernames and passwords through the smbusers and smbpasswd files in the /etc/samba directory.

There are drawbacks to the Samba Server Configuration utility. For example, you can't use it to edit all global parameters or share printers. You can't use it to set a Samba member server to join a domain.

Exam Watch

The Samba Server Configuration utility may not do everything you need. To configure most global settings, special printer shares, to join a domain, to control Samba services, your fastest option is to work from the command line interface.

Global Settings

To see what the Samba Server Configuration utility can do to the global settings in the smb.conf configuration file, choose Preferences | Server Settings. As you can probably guess from Figure 10-7, the basic settings set the **workgroup** and server **string** directives.

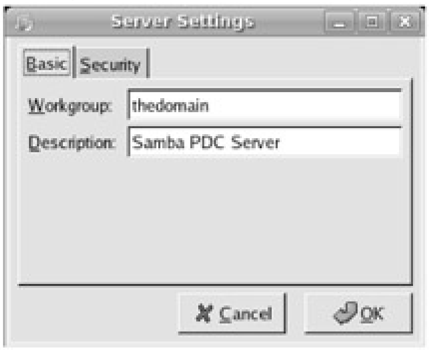

Figure 10-7: Samba Server basic settings

When you use this utility and assign default variables, it erases the variable from your smb.conf file. For example, if you set the **workgroup** name to **WORKGROUP**, this utility erases the **workgroup** command line from smb.conf. Therefore, it's an excellent idea to back up smb.conf before using the Samba Server Configuration utility.

In contrast, the Security tab supports a few more settings, as you can see in Figure 10-8. The entries are fairly straightforward. If you want more information on these variables, refer to the discussion on smb.conf earlier in this chapter:

- Authentication Mode sets the **security** value in /etc/samba/smb.conf. The default is **user**.
- Authentication Server sets up the location of the **password server**. There is no default.
- The Kerberos Realm is associated with an Active Directory user/password database and can be assigned only if **security = ads**.
- Encrypt Passwords is associated with the variable of the same name. The default is **yes**.
- Guest Account is associated with the variable of the same name. The default is **nobody**.

Figure 10-8: Samba Server security settings

On the Job

If you've selected a default, you may still see the variable in the smb.conf file in comments. Alternatively, you may see the variable in an unexpected location relative to the default comments.

Share Settings

Click the Add Share button. This opens the Create Samba Share window shown in Figure 10-9. The Basic tab helps you define the basic parameters associated with the share:

- Directory defines the directory that you want to share, using the path variable.
- By default, the Share Name is taken from the last part of the directory name. For example, if you're sharing the /usr/share/to/path1, Samba designates **[path1]** as the share name. You can also assign your own share name.
- Description allows you to define the comment that users can see in the browse list.
- Basic permissions lets you set **writable** as yes or no; read only (**writable=no**) is the default. (This is another case where two spellings for the same variable are acceptable; *writeable* is also an acceptable spelling in smb.conf.) The **visible** option, if activated, makes the share browsable.

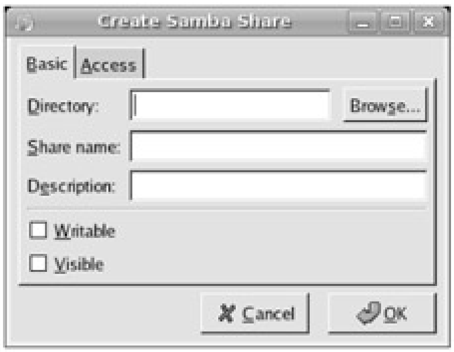

Figure 10-9: Basic components of Create Samba Share

The Access tab is simpler; it allows you to limit access to specific users from the smbpasswd configuration file. In other words, you can only limit access to users from the Samba password database. Once you've clicked OK, the Samba Server Configuration tool automatically updates the smb.conf configuration file.

Samba Users

The Samba Server Configuration tool also allows you to configure Samba users, based on the users already present in your /etc/passwd configuration file. Unfortunately, it can use only local password databases as of this writing. However, that's good enough to configure Microsoft usernames on this computer. To add Samba users from the Samba Server Configuration tool, choose Preferences | Samba Users. This opens the Samba Users window shown in Figure 10-10.

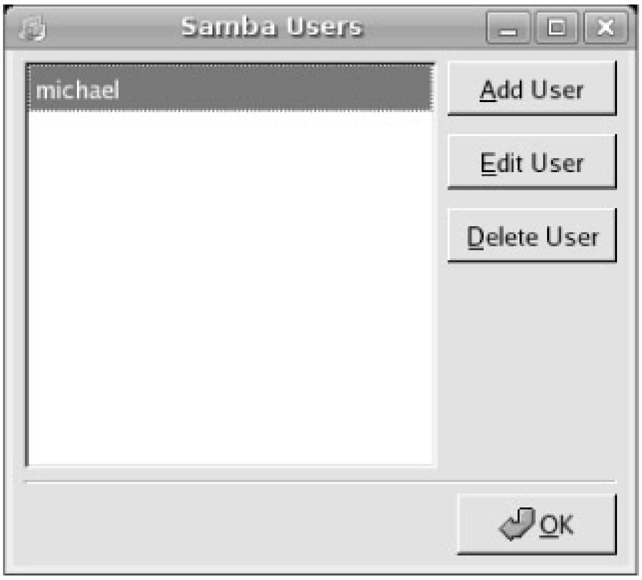

Figure 10-10: Current Samba users

As you can see, this window includes a list of currently configured Samba users. Click Add User. This opens the Create New Samba User window shown in Figure 10-11, where you can:

- Select an existing username from /etc/passwd.
- Enter the corresponding Microsoft Windows username.
- Set up a password for that Samba user. It can be different from that user's Linux password.

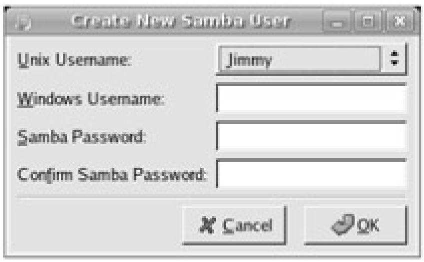

Figure 10-11: Creating a New Samba User

Click OK when you're done. Naturally, you can also change the Windows username and password for each Samba user, or even delete Samba users with the Edit User and Delete User buttons in the Samba Users window. Click OK to exit from the Samba Users window.

Creating a Public Share

Now you can create a public access share for use with the entire network. For the purpose of this chapter, create the /home/PublicShare directory. From the mail Samba Server Configuration screen, click Add Share to open the Create Samba Share window.

Enter the directory that you want to share, **/home/PublicShare**, in the Directory text box. Enter an appropriate Description, and select Writable and Visible. In the Access tab, select the Allow Access To Everyone option. Click OK, and exit from the Samba Server Configuration tool with the File | Quit command.

Now you'll have to finish the task directly from the text editor. The instructions so far add the following commands in the /etc/samba/smb.conf configuration file:

```
[
[PublicShare]

     comment = Shared Public Directory

      path = /home/PublicShare
```

Certification Summary

Networking services are an integral part of Red Hat Enterprise Linux. NFS, vsFTP, and Samba are a few of the services that you can configure for this operating system.

NFS allows you to share filesystems between Linux and Unix computers. This is a powerful method of controlling data and distributing I/O load, but there are many security concerns involved with its use. Be careful when setting up an NFS share on an unprotected network.

Red Hat includes one FTP server, the very secure FTP service. You can configure it in detail through the /etc/vsftpd/vsftpd.conf configuration file.

Samba allows a Linux computer to appear like any other Microsoft computer on a Microsoft Windows-based network. Samba is based on the Server Message Block protocol, which allows Microsoft computers to communicate on a TCP/IP network. It has evolved as Microsoft has adapted SMB to the Common Internet File System.

The main Samba configuration file, /etc/samba/smb.conf, includes separate sections for global settings and share definitions. The Red Hat Samba Server Configuration tool is a GUI tool that makes it easier to configure smb.conf. Changes to smb.conf can be easily tested with the **testparm** utility.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 10.

Configuring a Network File System (NFS) Server

? NFS is the standard for sharing files and printers between Linux and Unix computers.

? Key NFS processes are **rpc.mountd** for mount requests, **rpc.rquotad** for quota requests, and **nfsd** for each network share.

? NFS shares are configured in /etc/exports and activated with the **exportfs -a** command.

Client-Side NFS

? Clients can make permanent connections for NFS shares through /etc/fstab.

? If an NFS server fails, it can "hang" an NFS client. When possible, avoid using NFS on mission-critical computers.

? NFS and portmap have security problems. Limit their use when possible to secure internal networks protected by an appropriate firewall.

The File Transfer Protocol and vsFTPd

? RHEL includes the vsFTP server. The default configuration allows anonymous and real user access.

? You can customize vsFTP through the /etc/vsftpd/vsftpd.conf configuration file. It also uses authentication files in the /etc/vsftpd/ directory: ftpusers, user_list, and chroot_list.

Samba Services

? Samba allows Microsoft Windows computers to share files and printers across networks, using the Server Message Block (SMB) protocol on the TCP/IP protocol stack.

? Samba includes a client and a server. Variations on the **mount -t cifs** or **/sbin/mount.cifs** commands allow you to connect to a Microsoft Windows shared directory.

? The main Samba configuration file is /etc/samba/smb.conf. You can configure it in a text editor or a GUI tool such as the Samba Server Configuration tool.

? Samba allows you to configure your Linux computer as a member of a Microsoft Windows 9*x*-style workgroup.

? Samba allows you to configure your Linux computer as a Microsoft Windows server. It can also provide Microsoft browsing, WINS, and Domain Controller services, even on an Active Directory network.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

Configuring a Network File System (NFS) Server

1. In the /etc/exports file, if you want to export the /data directory as read-only to all hosts and grant read and write permission to the host superv in domain.com, what directive would you enter in that file?

____

2. Once you've configured /etc/exports, what command exports these shares?

____

3. Your company has just suffered an external security breach. As a result, the security group in your department has tightened the screws on all the servers, routers, and firewalls (but not SELinux). Up until this point, all user data had been mounted over NFS, but now nothing works. What's the most likely cause?

____

Answers

1. The following entry in /etc/exports would export the /data directory as read-only to all hosts and grant read and write permission to the host superv in domain.com:

```
/
/data (ro,sync)   superv.domain.com(rw,sync)
```

2. Once you've revised /etc/exports, the **exportfs -a** command exports all filesystems. Yes, you can re-export filesystems with the **exportfs-r** command. But there's no indication that NFS shares have yet been exported.

3. The most likely cause of NFS problems after security is boosted is an overzealous firewall.

Client-Side NFS

4. You're experiencing problems with NFS clients for various reasons, including frequent downtime on the NFS server and network outages between NFS clients and servers. What type of mounting can prevent problems on NFS clients?

____

Answers

4. Soft mounting can prevent problems such as lockups with NFS clients.

The File Transfer Protocol and vsFTPd

5. What default directive in /etc/vsftpd/vsftpd.conf should you disable if you don't want users logging into their accounts through the vsFTP server?

____

6. What directive should you enable if you want to keep regular users from getting to the top-level root directory (/) on your computer?

____

7. Based on the default RHEL 5 configuration, what file includes a list of users not allowed to log into the vsFTP server?

____

8. What additional directives do you need to add to the default vsFTP configuration file to allow security using PAM and TCP wrappers?

____

Answers

5. The default directive in /etc/vsftpd/vsftpd.conf that you should disable if you don't want users logging into their accounts through the vsFTP server is **local_enable=YES**.

6. The directive you should add if you want to keep regular users from getting to the top-level root directory (/) on your computer is **chroot_user=YES**.

7. Based on the default RHEL 5 configuration, both ftpusers and user_list in the /etc/vsftpd directory include a list of users not allowed to log into the vsFTP server.

8. The additional directives you need to add to the default vsFTP configuration file to allow security using PAM and TCP wrappers are **pam_service_name=vsftpd** and **tcp_wrappers=YES**.

Samba Services

9. A group that prefers Microsoft servers has set up a Windows 2000 server to handle file and print sharing services. This server correctly refers to a WINS server on 192.168.55.3 for name resolution and configures all user logins through the PDC on 192.168.55.8. If you're configuring the local Linux system as a PDC, what directive, at minimum, do you have to configure in the local Samba configuration file?

____

10. What command can be used to mount remotely shared Microsoft directories?

____

11. You made a couple of quick changes to your Samba configuration file and you need to test it quickly for syntax errors. What command tests smb.conf for syntax errors?

____

12. You've recently revised the Samba configuration file and do not want to disconnect any current users. What command forces the Samba service to reread the configuration file-without having to disconnect your Microsoft users or restarting the service?

____

Answers

9. At minimum, to configure a Linux system as a PDC, you need to configure the **security = user** directive.

10. The command that can be used to mount remotely shared Microsoft directories is **mount.cifs**.

11. The command that tests smb.conf for syntax errors is **testparm**.

12. The command that forces the Samba service to reread the configuration file-without having to disconnect your Microsoft users or restarting the service-is **service smb reload**.

PREV NEXT

Lab Questions

Lab 1

1. You'll need two Linux computers for this lab: one as an NFS server, and a second as an NFS client. Call these computers nfssvr.example.com and nfsclient.example.com. On the server, share the /home directories and provide write permissions to the client computer. On the client, set up the /home directory from the NFS server to be mounted the next time you boot that client computer.

Answers

1. This lab is the first step toward creating a single /home directory for your network. Once you get it working on a single client/server combination, you can set it up on all clients and servers. You can then use the NIS server described in Chapter 6 for a single Linux/Unix database of usernames and passwords for your network. On the NFS server, take the following steps:

    1. Set up some users and special files that you'll remember in some of the users' home directories on the server. The details are not important-just make a note of what you've done.
    2. Share the /home directory in /etc/exports with the nfsclient.example.com client. You can do this in this file with the following command:

Lab 2

2. Configure an FTP server for your computer. Make sure to allow only anonymous access. Don't allow anonymous users to upload to your server. Enable messages when users access your /var/ftp and /var/ftp/pub directories. Add an appropriate one-line message to each directory. Test the result, preferably from a remote computer. Start the vsFTP server and see that it starts automatically the next time you reboot your computer.

Answers

2. The vsFTP server is part of a simple package group. So if you have not installed this server during the installation process, the quickest thing to do is to connect to your installation source (CD or network) and install it from that location. For example, if the source is mounted on /mnt/source, you'd install it with the following command:

```
#
# rpm --Uvh /mnt/source/Server/vsftpd-*
```

Lab 3: Configuring Samba

3. This is a multi-part lab.

Part 1: Installing and Starting Samba

1. Ensure that all four components of the Samba service are correctly installed. What RPMs did you install and how did you install them?
2. Use one of the available service management tools to ensure that the Samba services are configured to start correctly when you boot Linux. What tool did you use?
3. Start Samba services now. You can use either the service management script located directly in /etc/rc.d/init.d or the "service" startup tool. How did you start your Samba service?
4. Verify that Samba services are running. How did you do this?

Part 2: Configuring Samba's Global Settings

1. You'll use Red Hat's Samba Server Configuration tool to configure your Samba service. Start this tool. If you didn't log in as the root user, did something happen before the tool started?
2. Configure the Samba global settings. You will provide workgroup services to your users. Set the workgroup name to something appropriate for your company.
3. Can you limit access to your company's domain name (such as example.com) through this tool? If you have to edit the Samba configuration file directly, what do you have to do?
4. Can you prevent access to evil.cracker.com through this tool? If you have to edit the Samba configuration file directly, what do you have to do?
5. Commit your changes. What do you need to do to make Samba reread the configuration file?

Part 3: Configuring File Shares

1. Open the main Samba configuration file.
2. Navigate to the predefined **[homes]** share.
3. Ensure that the **[homes]** share is available only to hosts on your example.com network.
4. Ensure that the share is writable to authenticated users but not available to guest users.
5. Commit your changes.

Answers

3. The chapter lab on Samba is designed to be easy to follow. However, you'll need explicit Linux knowledge to complete some specific steps. Answers to these steps can be found in the following:

Part 1

1. You've installed the Windows File Manager package group, which includes the samba-client, samba, and system-config-samba RPMs. These RPMs depend on the samba-common RPM, which you'll also need to install.
2. You can use the **chkconfig smb on** command or the Service Configuration utility described in Chapter 3 to make sure Samba starts the next time you boot Linux.
3. Use the **service smb start** command to begin the Samba service.
4. One way to verify Samba is to look for the existence of the **smbd** and **nmbd** processes in the process table. Use **ps aux | grep mbd** to see if these processes are present. Another way is with a service command such as **service smb status**.

Part 2

1. To use the Samba Server Configuration tool, you'll need the root password.
2. Many administrators stick with the standard Microsoft Windows **workgroup** name of **WORKGROUP** . You can find it in the output from the **smbclient -L //*clientname*** command.
3. If you want to limit access to your Samba server, you can't do it through the Samba Server Configuration tool. Set up the **hosts allow** command in /etc/samba/smb.conf.
4. If you want to restrict access from a specific computer to your Samba server, you can't do it through the Samba Server Configuration tool. Set up the **hosts deny = evil crackers.com** command in /etc/samba/smb.conf.
5. When you exit the Samba Server Configuration tool or save the smb.conf file, you can make Samba read the changes with the **service smb reload** command. But before committing the changes, you should test them with the **testparm** command.

Part 3

1. Open the main Samba configuration file, /etc/samba/smb.conf, in a text editor.
2. Navigate to the **[homes]** share in the last part of this file.

   Unless there is a limitation in the [global] section in this file, you can limit the **[homes]** share with the **hosts allow = example.com**. Commit your changes. Restart or reload the Samba daemon, **smb**, under the Status menu or with the appropriate **service** command.

Chapter 11: **Domain Name Service**

Overview

This chapter focuses on the Domain Name System (DNS), a service that translates human-readable domain names such as www.mommabears.com to IP addresses such as 10.245.43.5, and vice versa. DNS is a distributed database; each server has its own delegated zone of authority for one or more domains.

The DNS service associated with RHEL is the Berkeley Internet Name Domain (BIND). In this chapter, you'll learn how to edit and modify BIND configuration files to create authoritative DNS servers as well as slave and caching servers.

Once configured, there are a number of BIND utilities that can help find the systems on the local network as well as those on any other connected network, including the Internet.

On the Job

If you're interested in Dynamic DNS and Linux, one place to start is the Secure Dynamic DNS HOWTO from the Internet Engineering Task Force, at http://ops.ietf.org/dns/dynupd/secure-ddns-howto.html.

Inside the Exam

More Network Services

Both Red Hat exams require that you configure a Linux workstation as a client on a network. On a network with Linux computers, that naturally includes using DNS clients and servers.

This chapter may be related to two different items on the Troubleshooting and System Maintenance portion of the RHCT exam:

- Diagnose and correct misconfigured networking.
- Diagnose and correct host name resolution problems.

It's also related to the following item on the Troubleshooting and System Maintenance portion of the RHCE exam:

- Diagnose and correct problems with network services.

This naturally includes DNS. As noted in the Installation and Configuration requirements associated with the RHCE exam, this means being able to

- Configure DNS as a caching and a slave name service.

The omission of DNS as a master server in the Red Hat Exam Prep requirements is conspicuous by its absence. As always, Red Hat can change its requirements at any time (even halfway through the RHEL 5 release), so check the

Red Hat Exam Prep Web site for the latest information (https://www.redhat.com/training/rhce/examprep.html).

But as the only SELinux targeted setting relates to the overwriting of master zone files, there are very few SELinux issues addressed in this chapter.

Nevertheless, this chapter includes basic instructions for configuring a master DNS server.

Certification Objective 11.01-Understanding DNS: Zones, Domains, and Delegation

DNS, the Domain Name System, maintains a database that can help your computer translate domain names such as www.redhat.com to IP addresses such as 209.132.177.50. As no individual DNS server is large enough to keep a database for the entire Internet, each server is configured by default to refer requests to other DNS servers.

Basic Parameters

DNS on RHEL 5 is based on the **named** daemon, which is built on the BIND package developed through the Internet Software Consortium. (More information is available from the BIND home page at www.isc.org/products/BIND.) RHEL 5 includes BIND version 9.3. While this version of BIND supports the use of the /usr/sbin/rndc configuration interface, RHEL 5 still includes sample files based on the older /etc/ named.conf configuration file. However, you can use the rndc command to manage DNS operation, in the same way that you used apachectl to manage the Apache server.

Packages

If you're configuring your Linux computer solely as a DNS client, you can skip this section. The basic DNS client configuration files are automatically installed with even a minimal installation of RHEL 5.

On the other hand, if you're configuring your Linux computer as a DNS server, you'll need to install the packages associated with the DNS Name Server package group. You can do so with the package management tools described in Chapter 5. However, there are nine RPM packages associated with DNS:

- **bind** Includes the basic name server software, including /usr/sbin/named.
- **bind-chroot** Includes directories that isolate BIND in a so-called "chroot jail," which limits access if DNS is compromised.
- **bind-devel** Includes development libraries for BIND.
- **bind-libbind-devel** Contains the libbind BIND resolver library.
- **bind-libs** Adds library files used by the bind and bind-utils RPMs.
- **bind-sdb** Supports alternative databases, such as LDAP. Per the Red Hat Exam Prep guide and course outlines, I see no evidence that such relationships are covered on the Red Hat exams.
- **bind-utils** Contains tools such as **dig** and **host** that provide information about a specific Internet host. It should already be installed in any minimum installation of RHEL.
-

caching-nameserver Includes files associated with a caching nameserver.

- **system-config-bind** A GUI configuration tool useful for adding host and reverse address lookup data. It's not officially a part of the DNS Name Server package group.

These tools are easy to install from any Red Hat network installation source that you may have created in Chapter 2. Different options and commands for installing RPMs from a remote installation source are described in Chapter 5.

A DNS Client

There are two client configuration files associated with DNS: /etc/hosts and /etc/ resolv.conf. They are fairly straightforward, as described in the next section.

When your computer looks for another computer on a TCP/IP network such as the Internet, it typically looks in two places: /etc/hosts and any DNS servers that you've set up for your network. The order is determined by a single line in /etc/ nsswitch.conf:

```
hosts : files ,dns
```

Certification Objective 11.02-The Berkeley Internet Name Domain (BIND)

You can configure a DNS server by directly editing the DNS configuration files. Alternatively, you can configure a DNS server using the Red Hat Domain Name Service configuration tool. Careful use of both tools can help you learn more about DNS.

You can set up four different types of DNS servers:

- A master DNS server for your domain(s), which stores authoritative records for your domain.
- A slave DNS server, which relies on a master DNS server for data.
- A caching-only DNS server, which stores recent requests like a proxy server. It otherwise refers to other DNS servers.
- A forwarding-only DNS server, which refers all requests to other DNS servers.

On the Job

system-config-bind is the successor to *bindconf*. Red Hat Enterprise Linux includes a link from *bindconf* to *system-config-bind*.

The DNS Configuration Files

DNS configuration files are required to configure your Linux computer as a client and as a server. The client configuration was described earlier in this section. There are a number of additional configuration files that support the use of DNS as a server, as described in Table 11-1. You may have to create some of these files as you configure DNS. If you've installed the bind-chroot package, you'll find the real copy of these files in the /var/named/chroot/var/named directory, linked to /var/named.

Table 11-1: DNS Server Configuration Files

DNS Configuration File	Description
/etc/sysconfig/named	Set up different configuration and data file directories through this file.
/etc/named.conf	The main DNS configuration file. Incorporates data from other files with the **include** directive.
/etc/named-caching-nameserver.conf	A template DNS configuration file for a caching DNS server.
/etc/named.rfc1912.zones	Adds appropriate zones for localhost names and addresses.

/var/named/chroot/etc/rndc.key	The authentication key required to support requests to the DNS server.
/var/named/my.internal.zone.db	The zone file for the local network.
/var/named/slaves/my.slave.internal.zone.db	The zone file for a slave DNS server.
/var/named/slaves/my.ddns.internal.zone.db	The zone file for a dynamic DNS server.
/var/named/localdomain.zone	The zone file for the localhost's domain.
/var/named/localhost.zone	The zone file for the localhost computer.
/var/named/named.broadcast	A broadcast record for the localhost.
/var/named/named.ca	A list of root DNS servers on the Internet.
/var/named/named.local	A reverse zone record for the localhost.
/var/named/named.ip6.local	An IPv6 version of named.local.
/var/named/named.zero	Defaults to the broadcast record for the localhost.
/var/named/data/named.stats.txt	Statistics from your DNS server, only available after DNS is active.

As noted in the default copy of /etc/sysconfig/named, the **ROOTDIR** directive points to /var/named/chroot/ as the root directory for all other DNS server configuration files. We assume that you've installed bind-chroot, in which case, all files in Table 11-1 are in that directory unless otherwise noted. Some may be linked to that directory. For example, once created, /etc/named.conf is actually a soft link to /var/ named/chroot/etc/named.conf.

In many cases, you do not have to use the file names listed in Table 11-1. All you need to do is make sure the file is properly cited in the main named.conf configuration file. Remember that if you've installed the bind-chroot RPM, these files will be linked to files in the /var/named/chroot/ directory.

Many of these files include templates in the DNS documentation directory, /usr/share/doc/bind-*versionnum*/sample.

In the following sections, you'll see how to configure the files that you need for a working DNS server. But first, you should know how to configure your computer as a DNS client. One thing to remember is that all of the files in /var/named include a dot at the end of each domain name. For example, /var/named/named.local lists the local computer as:

```
l
bca lbs t.
```

Certification Objective 11.03-BIND Utilities

Once you've configured BIND as a DNS server, there are a number of commands you can use to keep it working. Red Hat even has its own GUI configuration tool for BIND. While I describe the Red Hat tool briefly, it's somewhat complex. Red Hat didn't even include a GUI tool for RHEL 4, so its reliability in RHEL 5 may not have been fully tested in production.

BIND Commands

There are three commands associated with the BIND service: rndc, **host**, and **dig**. The rndc command is a better way to control the service. The rndc and **host** commands are successors to **nslookup**.

The rndc commands are straightforward. When you run rndc by itself, the output guides you through the available options. The options I use are straightforward: rndc **stop** and **rndc start** don't require explanation. The **rndc reload** command rereads any changes you've made to the configuration or DNS database files. The **rndc status** command confirms that DNS is running, along with information on the DNS database.

While you can still use commands such as **service named start** and **service named reload**, the rndc command can do more. Because the current Red Hat Exam Prep guide suggests that you need only know how to create a caching and slave nameserver, the details are not important for this book.

After you configure DNS and make it reread your configuration files with the **rndc reload** command, examine the results with the **host -l example.com** command. I've shown the results from my zone file in Figure 11-7.

```
[root@Enterprise5vm ~]# host -l example.org
example.org name server enterprise5vm.
drakelaptop.example.org has address 192.168.0.17
drakeoffice.example.org has address 192.168.0.13
enterprise5a.example.org has address 192.168.0.50
Enterprise5vm.example.org has address 192.168.0.15
[root@Enterprise5vm ~]# █
```

Figure 11-7: Listing a working DNS zone

Now test the setup. Use the **dig** command to examine your work. For example, if you use **dig** to look up the address of www.redhat.com, you'll see something like the output shown in Figure 11-8.

```
[root@Enterprise5vm ~]# dig www.redhat.com

; <<>> DiG 9.3.2 <<>> www.redhat.com
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: SERVFAIL, id: 14785
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.redhat.com.                IN      A

;; Query time: 101 msec
;; SERVER: 192.168.0.15#53(192.168.0.15)
;; WHEN: Tue Jul 25 12:35:51 2006
;; MSG SIZE  rcvd: 32

[root@Enterprise5vm ~]# █
```

Figure 11-8: DNS query using dig

The dig command asks your DNS server to look for the www.redhat.com server. Assuming IP address information for www.redhat.com isn't stored locally, it then contacts one of the name server computers listed in /etc/resolv.conf. If that doesn't work, it goes to one of the name servers listed in the named.ca file and makes its requests from there. The request may be passed on to other DNS servers. Therefore, it can take some time before you see an answer.

On the Job

The *nslookup* command is now deprecated, and you may not even be able to use it in a future release of Red Hat Enterprise Linux.

The DNS Configuration Tool

Red Hat has created a number of excellent GUI configuration tools. They are "front ends" that can help many administrators create the configuration files that they need. While the Red Hat Domain Name Service tool is promising, it was just introduced for RHEL 5. Red Hat did not even include a GUI configuration tool for BIND in RHEL 4.

In any case, it is best if you learn how to configure Linux services, including DNS, directly from the configuration files. As a Linux systems administrator, you may not always have access to the GUI. You may need to administer servers remotely, which makes GUI configuration difficult at best.

If you want to try the Red Hat Domain Name Service configuration tool, back up your DNS configuration files first: /etc/named.conf, as well as the files in the /var/ named directory (subdirectories, and links actual files in other directories).

Exercise 11-1: Setting Up Your Own DNS Server

Following the example files shown previously, set up your own DNS server. Set it up to serve the domain called rhce.test. As long as your domain is private, it doesn't matter that rhce.test does not match the standard domain name types such as .com or .net.

1. Edit the /etc/named.conf file to reflect the configuration files that you plan to use. Name the zone file rhce.test.zone and set it to be a master domain.
2. Edit the file /var/named/rhce.test.zone and place the proper zone information in it. Start by adding in the header with the serial number and expiration information.
3. Add the SOA resource record (RR) with a proper administrative e-mail address contact.
4. Add NS and MX RRs for the domain. Use the 192.168.0.0/24 address range. If you're configuring an actual TCP/IP network with static IP addresses, feel free to use the assigned IP addresses on your network.
5. Add several hosts to the zone file. Use WWW, FTP, and mail for a few.
6. Save the zone file and then restart **named** with the **rndc reload** command.
7. Use the **dig** command to check the rhce.test domain. If it works, you have a working DNS server.

Certification Summary

DNS provides a database of domain names and IP addresses that help Web browsers and more find sites on the Internet. The default DNS server for RHEL uses the Berkeley Internet Name Domain (BIND). It's a distributed database for which each administrator is responsible for his or her own zone of authority.

DNS can be controlled with various rndc commands. The diagnostic tool for DNS is now **dig**. You can also use the more traditional **nslookup** command, but it has been deprecated in RHEL 5.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 11.

Understanding DNS: Zones, Domains, and Delegation

? DNS is based on the Berkeley Internet Name Domain (BIND), using the **named** daemon.

? Key packages include bind-chroot, which adds security by supporting DNS in a chroot jail.

The Berkeley Internet Name Domain (BIND)

? Critical DNS configuration files include /etc/named.conf and the files in the /var/named directory.

? Caching-only DNS servers store requests and their associated IP addresses on a computer.

? Slave DNS servers need to point to a master DNS server, with the appropriate **masters** directive in /etc/named.conf.

? Every time you change DNS, remember to update the serial number in your zone file. Otherwise, other DNS servers don't realize that you've changed anything.

BIND Utilities

? There are a number of BIND utilities that can help you manage the service, including rndc and **rndc-confgen**. Others can help you check the database, including **dig** and **host**.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

Understanding DNS: Zones, Domains, and Delegation

1. If you're configuring a connection from a client to a DNS server, what file would you use?

2. If your ISP has a DNS server address of 10.11.12.13, what directive would you add to the DNS client configuration file used in question 1?

Answers

1. If you want to configure a connection from a client to a DNS server, you would use /etc/resolv.conf.

2. If your ISP has a DNS server address of 10.11.12.13, you would add the following directive to the /etc/resolv.conf file:

```
nameserver 10.11.12.13
```

The Berkeley Internet Name Domain (BIND)

3. If you configure DNS communication on port 53, what changes would you make to a firewall to support access by other clients to the local DNS server?

4. What file includes a basic template for a DNS caching nameserver?

5. What file includes a basic template for a master DNS server?

6. If you've installed the bind-chroot RPM, where will you find the actual DNS server configuration files?

7. If you've installed the bind-chroot RPM, where will you find links to the actual DNS server configuration files?

____

8. Why would you configure a reverse DNS zone database?

____

9. If there are errors when you start the DNS service, where are error messages available by default?

____

10. What command makes sure that the DNS service starts the next time you boot Linux?

____

Answers

3. To support access by other clients to the local DNS server, make sure TCP and UDP traffic is supported through the firewall on port 53.

4. The /etc/named.caching-nameserver.conf file includes a basic template for a DNS caching nameserver. Alternatively, named.conf in the /usr/share/system-config-bind/profiles/default directory also includes a default caching nameserver template.

5. The /etc/named.rfc1912.zones file includes a basic template for a master DNS server. Sample files are also available in the /usr/share/doc/bind-*versionnum*/sample/ directory.

6. If you've installed the bind-chroot RPM, the actual DNS server configuration files can be found in the /var/named/chroot/etc and /var/named/chroot/var/named directories.

7. If you've installed the bind-chroot RPM, you can find links to the actual DNS server configuration files in the /etc and /var/named directories.

8. You would configure a reverse DNS zone database to help other services such as sendmail verify the domain names associated with IP addresses.

9. If there are errors when you start the DNS service, error messages are available by default in /var/log/messages.

10. The command that makes sure that the DNS service starts the next time you boot Linux is

```
#
# c kco nfig ame do n
```

BIND Utilities

11. What command lists the data associated with the example.net domain on a properly configured DNS server?

____

12. After you revise the DNS database files, what command most appropriately rereads the database, without restarting the service?

____

Answers

11. The command that lists the data associated with the example.net domain on a properly configured DNS server is

```
#
# host -l example.net
```

12. After you revise the DNS database files, the command that most appropriately rereads the database, without restarting the service, is

```
#
# rndc reload
```

PREV NEXT

Lab Questions

Lab 1

1. In this lab, you'll set up a caching DNS nameserver on your network. Use the /etc/named.caching-nameserver.conf file, modify appropriate files on the clients on your network, and make sure the appropriate daemon is active and starts the next time you boot Linux.

Answers

1. In this lab, you have the benefit of the /etc/named.caching-nameserver.conf configuration file. All you need to do is:

    1. Copy the template configuration file to /etc/named.conf.

    2. Modify the **listen-on port 53** directive to include the local IP address; for example, if your IP address is 10.11.12.13, the directive will look like:

Lab 2

2. Your internal network is growing, and you're having trouble keeping up with the different workstations that are being added on a regular basis. You use the good.example.com subdomain for your internal network, and you've named your computers for your departments, such as engr1 through engr10.good .example.com.

Your mail server is named postal, your Web server is named www, your FTP server is named ftp. You want to configure a DNS server on the computer named names. What do you need to do?

While you may not have enough information in this lab to create a complete and working file, you should be able to determine an outline of what you need to do, with the possible exception of specific IP addresses.

Answers

2. While you could subcontract out the task to an ISP, it's easy to create a DNS server for your internal network. The basic files are already available in RHEL 5. All you need to do is modify these files and add appropriate zone files to your /var/named/chroot/var/named directory. I'll describe the basics on how you can set up a DNS server by directly editing the appropriate configuration files. Assume that you're using the 10.11.12.0/255.255.255.0 network addresses for your LAN.

First, you'll need to modify the default /etc/named.conf configuration file. It's best to start by backing up this file. You'll need to add stanzas that refer to a zone and a reverse zone file. The stanzas are straightforward:

```
zone "goodexample.com" IN{
      type master;
      file "goodexample.com.zone";
}
};
```

Chapter 12: Electronic Mail

Overview

Linux offers a number of alternative methods for handling incoming and outgoing e-mail. Red Hat Enterprise Linux includes sendmail, Postfix, and Dovecot for this purpose. (Yes, it includes exim and squirrelmail as well, but sendmail, Postfix, and Dovecot are what's listed in the RH300 course outline, and Dovecot is the new default for incoming e-mail.)

Perhaps the most common server for outgoing e-mail is sendmail, which may already be installed on your RHEL system. Once it is installed and configured, you can set up sendmail as your own personal mail server (subject to the limitations of your ISP). RHEL includes the open-source version of sendmail; the commercial version is known as Sendmail (with the capital *S*). One alternative to sendmail that is installed on RHEL is known as Postfix.

RHEL includes Dovecot for standard incoming e-mail protocols. It's relatively easy to configure these protocols, including POP3 (Post Office Protocol), POP3S (the secure version), IMAP (Internet Message Access Protocol), and IMAPS (the secure version).

Exam Watch

A number of alternatives to sendmail are not covered in this book; they include procmail, mail.local, exim, Cyrus IMAP, and uucp. Only sendmail, Postfix, and Dovecot are currently part of Red Hat's public RH300 course outline.

Inside the Exam

More Network Services

This chapter is focused on the RHCE exam. While the ability to configure an e-mail client is a prerequisite skill for both exams, the Red Hat Exam Prep guide suggests that RHCE candidates must know how to configure the following network services:

- SMTP
- IMAP, IMAPS, and POP3

While sendmail and Postfix are not specified in the Exam Prep guide, the RHCE course, RH300, includes coverage of Dovecot, sendmail, and Postfix. Dovecot is one RHEL 5 service that can handle IMAP, IMAPS, and POP3.

This chapter is related to the following items on the Troubleshooting and System Maintenance portion of the RHCE exam:

- Diagnose and correct problems with network services, including whatever mail related services that you've configured.

There's also the related Troubleshooting and System Maintenance requirement to

-

Diagnose and correct networking services problems where SELinux contexts are interfering with proper operation.

But the only e-mail related option is to disable SELinux protection for Postfix (and fetchmail).

You may need to install, configure, and secure these services during the Installation and Configuration portion of the exam. And don't forget to make sure any required services are active when you reboot, or you may not get credit for your work.

As for the RHCE exam, you may have to configure or troubleshoot the e-mail services discussed in this chapter. So as you read this chapter and look through the configuration files and exercises, be willing to experiment. And practice, practice, practice what you learn.

PREV NEXT

Certification Objective 12.01-Mail Transport Agents, Mail Delivery Agents, and Mail User Agents

When you install sendmail, Postfix, and/or Dovecot, you also get huge and difficult-to-read configuration files. Do not be intimidated, as it's likely that you'll have to change only a few entries in each file.

Definitions

A mail server has three major components, as described in Table 12-1. You need all three components to have a fully functional mail system. Fortunately, as the other components are already installed, you should have to install only the MTAs that you need on a standard RHEL system.

Table 12-1: Mail Server Components

Abbreviation	**Meaning**	**Examples**
MTA	Mail transfer agent	sendmail, Postfix, Dovecot
MUA	Mail user agent	mail, Evolution, elm
MDA	Mail delivery agent	procmail, maildrop

On any Linux computer, you can configure some mail transfer agents (sendmail or Postfix) for various outbound services, such as forwarding, relaying, method of transport (such as TCP or UDP), lists of computers with other MTAs, optional aliases, and spooling directories. Others, such as Dovecot, are designed to handle only incoming e-mail services.

E-mail systems are heavily dependent on name resolution. While you could handle name resolution through /etc/hosts on a small network, any mail system that requires Internet access needs access to a fully functional DNS server.

The sendmail and Postfix systems use SMTP to send e-mail. But that is only one end of the mail system. You also need to use a service such as Dovecot to enable POP3 and/or IMAP (or the secure cousins, POP3s and IMAPS) to receive e-mail.

On the Job

While this chapter refers to the IMAP and IMAPS protocols, rest assured that these options support the current versions of these protocols, IMAP4 and IMAP4S.

SMTP, the Simple Mail Transfer Protocol, has become one of the most important service protocols of the modern era. Much of the Internet-connected world lives and dies by e-mail and relies on SMTP to deliver it. Like POP3 and IMAP, SMTP is a *protocol*, a set of rules for transferring data used by various mail transfer agents.

Installing Mail Server Packages

The RPM packages associated with sendmail and Postfix are both part of the Mail Server package group. Key packages are listed in Table 12-2. You can install them with the **rpm** or yum command. Just remember that you may

not need to install everything in this table.

Table 12-2: Key Mail Server RPMs

RPM Package	Description
cyrus-imapd*	Installs the Cyrus IMAP enterprise e-mail system (several packages); may require perl-Cyrus
cyrus-sasl	Adds the Cyrus implementation of the Simple Authentication and Security Layer (SASL)
dovecot	Supports both the IMAP and the POP incoming e-mail protocols
exim	Adds another MTA; another alternative to sendmail and Postfix
mailman	Supports e-mail discussion lists
postfix	Includes an alternative to sendmail
sendmail	Installs the most popular mail server of the same name
sendmail-cf	Adds a number of templates that you can use to generate your sendmail configuration file
spamassassin	Includes the spam fighting package of the same name
squirrelmail	Installs a Web-based e-mail server
system-switch-mail system-switch-mail-gnome	Adds a GUI method for switching between sendmail and Postfix

When you install the default Mail Server package group, you're installing the sendmail and Dovecot packages. Since you may not need all of these packages, it may be faster to install these with the **rpm** or yum command, especially if you're configuring your Linux computer from the text console. It takes time to start the GUI.

On the Job

You can find a list of RPMs associated with each package group on the first installation CD in the /Server/repodata directory in the comps-rhel5-server-core.xml file. If you're running the RHEL 5 desktop, substitute "Client" for "Server."

Certification Objective 12.02-Reception with Dovecot

Once you've installed the Dovecot package, it's easy to configure. All you need to do is add the appropriate directive in the Dovecot configuration file and make sure it starts the next time you start Linux. If you use a secure incoming e-mail protocol such as POP3S or IMAPS, you'll want to configure an appropriate certificate as well.

The process has changed significantly relative to RHEL 4. But first, let's review the POP3 and IMAP protocols.

POP

The Post Office Protocol (POP) is one of the major mail delivery protocols. It includes some basic commands that allow you or an e-mail client to send and retrieve messages. A mail service can be configured to be a central depository for incoming mail messages from any other MTA service. Client applications then download the mail messages from the POP server for processing at the local host.

The current version of POP is known as POP3.

On the Job

You can configure user accounts that are designed to service only POP user accounts, where users log in and receive mail only and no interactive service is provided. Just set up the appropriate mail client in the login configuration sequence for a given user.

IMAP

The IMAP (Internet Message Access Protocol) is the other major mail delivery protocol. While POP downloads all e-mail to the client, an IMAP server maintains all mail messages on the server, as a database. IMAP is commonly used by businesses that service users who log in from different locations. It's also the most common mail delivery protocol for Web-based mail services.

The current version of IMAP is known as IMAP4.

Configuration File

Now that you've reviewed the protocols, let's start configuring Dovecot. The main configuration file, /etc/dovecot.conf, is well commented. As the file is nearly 1000 lines long, this section focuses on a few key directives. The first thing to note is this comment:

```
#
# Default values are shown after each value, it's not required to
#
```

Certification Objective 12.03-sendmail Configuration

The sendmail daemon is configured from a directory of files in /etc/mail and a directory of configuration files in /usr/share/sendmail-cf. There are two basic configuration files: sendmail.cf for incoming mail and submit.cf for outgoing mail. I describe the key configuration files in /etc/mail in a bit of detail here.

- **sendmail.cf** The main sendmail configuration file.
- **sendmail.mc** A macro that's easier to edit, which can be used to generate a new sendmail.cf file.
- **access** Supports outgoing access control to your sendmail server. The default version of this file supports access from the local computer. You can add host names or networks to this list, with a message to **REJECT** with an error message, **DISCARD** without an error message, or **RELAY** to accept and send the e-mail.
- **domaintable** Allows you to map different domains. For example, if you've changed your domain name from Osborne.com to Mcgraw-hill.com, people might still send e-mails to addresses such as michael@Osborne.com. The following line would forward that e-mail to michael@Mcgraw-hill.com.

Certification Objective 12.04-Configuring and Activating Postfix

You can configure the Postfix mail server as a substitute for sendmail. The configuration files are stored in the /etc/postfix directory. The main configuration file, main.cf, is somewhat simpler than sendmail.cf, as it includes over 600 lines. Back up this file and open it in a text editor. There are several things that you should configure in this file to get it working-and limit access to your system and network:

- Activate and modify the following **myhostname** directive to point to the name of your computer:

Certification Objective 12.05-Selecting an E-mail System

If you have installed and configured both sendmail and Postfix, be careful. Don't activate both. There are two simple commands that can help you select the default e-mail server: **alternatives** and **system-switch-mail**.

Once you select a system, you need to know how to test the system. Sure, you can set up Thunderbird or Evolution, but that takes time. It's easiest to use the **mail** utility to send a quick e-mail. When you do so from a remote system, you can make sure the active server does what you want, and does not accept mail from other systems.

Using alternatives to Select an E-mail System

The **alternatives** command is straightforward. Try the following:

```
#
# alternatives --config mta
```

Certification Summary

Red Hat uses the Dovecot service for incoming e-mail. It can be configured for use with both the POP3 and IMAP protocols, as well as their secure cousins (POP3S and IMAPS). If you want to handle secure incoming e-mail, you'll want to set up SSL certificates in the /etc/pki/dovecot/certs and /etc/pki/dovecot/private directories. The mkcert.sh script is designed to help create custom certificates, using settings you add to the /etc/pki/dovecot/dovecot-openssl.cnf file.

Red Hat includes two servers for outgoing e-mail: sendmail and Postfix. Both rely on SMTP to send e-mail. The sendmail service includes difficult-to-read configuration files: sendmail.cf and submit.cf, both in the /etc/mail directory. Fortunately, you can configure these files through easier-to-read macro files, sendmail.mc and submit.mc. Postfix is somewhat easier to configure directly through its own configuration file, /etc/postfix/main.cf. Red Hat also provides a variety of e-mail clients, including **mail**, mutt, and the mail clients that come with the various Web browsers that are also available for Linux. These clients use either POP or IMAP protocols to receive e-mail.

Because RHCE is a performance-based exam, just like life, it is important to practice all the skills discussed in this chapter. You may need to use these skills on the exam!

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 12.

Mail Transport Agents, Mail Delivery Agents, and Mail User Agents

? Mail transport agents (MTA) are servers that carry e-mail to and from users.

? Three MTAs include sendmail, Postfix, and Dovecot.

? Mail user agents (MUA) are e-mail clients such as mail and Evolution.

? Mail delivery agents (MDA) are used with MTAs to transmit e-mail.

Reception with Dovecot

? The Dovecot service supports receipt of both POP3 and IMAP4 e-mail.

? Dovecot supports secure SSL versions of POP3 and IMAP4; you'll need to create SSL certificates in appropriate /etc/pki/dovecot directories.

sendmail Configuration

? The main sendmail configuration file is /etc/mail/sendmail.cf. It's easier to configure sendmail through its macro file, /etc/mail/sendmail.mc.

? You can customize the computers allowed to access your sendmail server through the access and virtusertable files in the /etc/mail directory.

? The **make -C /etc/mail** command processes all files in that directory, to configure sendmail completely.

Configuring and Activating Postfix

? The Postfix server is easier to configure through configuration files in the /etc/postfix directory. In fact, you can configure the main.cf file directly.

? Postfix also uses /etc/aliases to map and forward e-mail from old to new addresses.

Selecting an E-mail System

?

RHEL 5 allows you to select between sendmail and Postfix. Don't activate both.

?

Available tools for selecting a service are **alternatives** and **system-switch-mail**.

?

Different mail clients can receive e-mail using the POP or IMAP protocols. Mail clients such as **mail** and mutt are available at the command line. GUI mail clients such as Thunderbird and Evolution are also available.

PREV NEXT

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of the questions.

Mail Transport Agents, Mail Delivery Agents, and Mail User Agents

1. What package described in this chapter can be configured for incoming POP3 and IMAP4 e-mail? ?

Answers

1. The package that can be configured to receive POP3 and IMAP4 e-mail is Dovecot.

Reception with Dovecot

2. If you don't want to allow secure incoming e-mail, what would you do with the following directive in /etc/dovecot.conf? ?

```
#
# protocols = imap imaps pop3 pop3s
```

3. What script should you use to configure an SSL certificate for Dovecot? ?

4. What certificate files do you need to move or delete before creating a customized SSL certificate? ?

Answers

2. If you don't want to allow secure incoming e-mail (for whatever reason), you need to activate and change the **protocols** directive to specify only the nonsecure e-mail protocols:

```
p
protocols = imap pop3
```

3. The script that helps configure an SSL certificate for Dovecot is mkcert.sh, in the /usr/share/doc/dovecot-1.0/examples directory.

4. The certificate files that you need to move or delete before creating a customized SSL certificate are both named dovecot.pem, in the /etc/pki/dovecot/certs and /etc/pki/dovecot/private directories.

sendmail Configuration

5. In what file would you store forwarding e-mail addresses?

6. Why would you want to comment out the following directive in sendmail.mc?

```
DAEMON_OPTIONS(`Port=smtp,Addr=127.0.0.1, Name=MTA') dnl
```

7. What do you need if you want to comment out the following directive in sendmail.mc?

```
FEATURE(`accept_unresolvable_domains') dnl
```

8. What command processes all files in /etc/mail?

Answers

[5.](#) Forwarding e-mail addresses for both sendmail and Postfix are normally stored in /etc/aliases. If you're forwarding e-mail for entire domains, the appropriate file is /etc/mail/domaintable. Make sure to process these files into appropriate databases; for /etc/aliases, the database is updated with the **newaliases** command. For /etc/mail/domaintable, the database is updated with the next **make -C /etc/mail** command.

[6.](#) The noted directive limits sendmail server access to the localhost computer, IP address 127.0.0.1.

[7.](#) If you don't want to accept unresolvable domains, say to minimize spam on your system, you'll need reliable DNS service. Otherwise, reverse DNS searches may fail, and your system may not accept even legitimate e-mail.

[8.](#) The **make -C /etc/mail** command processes all files in /etc/mail, including sendmail.mc, submit .mc, and the database files in that directory. It does not process any forwarding aliases in /etc/ aliases; if you make a change to this file, you need to use the **newaliases** command.

Configuring and Activating Postfix

9. How would you change the following directive in /etc/postfix/main.cf to open Postfix to all systems?

```
inet_interfaces = localhost
```

10. If you use /etc/aliases for forwarding e-mail, what command processes these files into an appropriate database file for Postfix?

Answers

[9.](#) The simplest solution is to change the directive to

```
inet_interfaces = all
```

10. Forwarding e-mail addresses for both sendmail and Postfix are normally stored in /etc/aliases. Make sure to process these files into appropriate databases; for /etc/aliases, the database is updated with the **newaliases** command.

Selecting an E-mail System

11. What command allows you to select between the Postfix and sendmail servers?

____

12. What command line e-mail client can you use to test a server?

____

Extra Credit: What can you run directly from the command line to send a simple test e-mail?

____

Answers

11. There are a wide variety of solutions to this problem. The simplest solutions are fastest and best. Two simple commands can help you switch between Postfix and sendmail: **alternatives --config mta** and **system-switch-mail**.

12. Several command-line e-mail clients are available for RHEL, including mutt and **mail**.

Extra Credit: One simple way to send a message with the **mail** client is with the following command (variations are acceptable, as the point is to check the e-mail server quickly from a remote location). The e-mail subject in this case is *test subject* and the message is *hello kitty*.

```
#
# echo "hello kitty" | mail -s "test subject" michael@enterprise5a
```

PREV
NEXT

Lab Questions

Lab 1

1. Configure Thunderbird or Evolution to read your e-mail. You can set it up to read e-mail without downloading it from the server (even if it's a POP3 server).

Run the **mail** command as the root user. Browse through available e-mails. Normally, you should have more than a few based on issues such as bad logins and logwatch jobs.

Review the appropriate mail server spool and log files. Current Postfix mail log messages can also be found in /var/log/maillog.

Answers

1. This should be a trivial lab for most users. Anyone who is preparing for the RHCE exam should already know how to configure GUI e-mail clients such as Thunderbird and Evolution. This part of the lab is designed to get you to think about what these e-mail clients do and how you can configure e-mail clients.

But there is more. The **mail** command opens the e-mail client of the same name. It's already linked to mails in the local account, collected in /var/spool/mail/*USERNAME*.

When you review log messages in /var/log/maillog, you may see some errors. For example, if you see something similar to this,

```
f
fatal: bind 0.0.0.0 port 25: Address already in use
```

Lab 2

2. Configure a Dovecot server with support for regular and secure POP3 and IMAP4 services. Create appropriate SSL certificates in default directories, based on your location.

Answers

2. This lab assumes you haven't protected access with a firewall, tcp_wrappers, or SELinux, as described in Chapter 15. One way to perform this lab is with the following steps:

1. Make sure the dovecot RPM is installed. Open /etc/dovecot.conf, and activate the following directive:

Lab 3

3. Set up a sendmail mail server for your network. First, make sure to disable local-only access in the /etc/mail/sendmail.mc file. Add your network to the /etc/mail/access file. Test the results, preferably from a remote computer on your network.

   Set up two users on your system. You can use existing users. For the purpose of this lab, assume these users are linus and bill. Make sure e-mail is forwarded from bill to linus.

   Make sure to start the sendmail server now, and see that it starts automatically the next time you reboot your computer.

Answers

3. In sendmail, to disable local-only access in the /etc/mail/sendmail.mc file, comment out the following line. Unlike most Linux configuration files, the comment code is a **dnl** at the start of this line:

```
DAEMON_OPTIONS(`Port=smtp,Addr=127.0.0.1, Name=MTA') dnl
```

Lab 4

4. Set up a Postfix mail server for your network. Test the results, preferably from a remote computer on your network. Retain the forwarding from Lab 2, where e-mail addressed to user bill is forwarded to user linus. Connect to a remote system, and send e-mail to user bill.

Answers

4. In Postfix, to disable local-only access in the /etc/postfix/main.cf file, change the **inet_interfaces** directive to accept **all** connections:

```
inet_interfaces = all
```

PREV NEXT

Chapter 13: Other Networking Services

Overview

This chapter starts with a description of the Extended Internet Services Daemon (xinetd), also known as the "Internet Super Server." It governs a lot of services that do not have their own individual daemons. The techniques required to configure and control xinetd services are subtly different from regular services.

This chapter continues with a discussion of the Secure Shell package, which allows you to connect remotely to the systems you need to administer. While it already encrypts what you type over a network, you can configure it with more levels of security. It's also one way to configure remote connections to GUI applications.

DHCP allows a Linux computer to serve dynamic IP addresses. You can configure a range of IP addresses, reserve a specific IP address for the hardware address associated with a client's network card, and assign more information such as the gateway and DNS IP address to every system that requests an IP address.

Finally, the Network Time Protocol (NTP) can help keep the computers on a network running on the same time. That can be important for logs and for keeping backup servers in sync, which can be especially important for financial transactions.

Certification Objective 13.01-The Extended Internet Services Daemon (xinetd)

Linux typically supports network communication between clients and servers. For example, you can use Telnet to connect to a remote system. The Telnet client on your computer makes a connection with a Telnet server daemon on the remote system. This section assumes that you've installed the default RHEL krb5-workstation RPM package, which includes a more secure version of Telnet.

While the focus in this section is on the Kerberos-secured Telnet, other xinetd packages of note include rsync, which is popular for backups; cvs, popular for software development version control; and gssftp, which is a Kerberos-secured FTP service. As no xinetd service is explicitly cited in the RHCE Exam Prep guide, I keep the coverage of xinetd services to a minimum.

Inside the Exam

More Network Services

Both Red Hat exams require that you configure a Linux workstation as a client on a network. However, the secure shell is not included in the current version of the RHCT portion of the Exam Prep guide. And client configuration with a DHCP server was covered in Chapter 7. So this chapter focuses on the needs of RHCE candidates.

The Red Hat Exam Prep guide suggests that you can expect to configure SSH services during the exam. It also suggests that you need to know how to configure host- and user-based security for each service. While much of that is discussed in Chapter 15, some security configuration options are available for SSH.

While neither DHCP servers nor xinetd services are listed in the Red Hat Exam Prep guide, they are listed in the RH300 course outline, the prep course for the RHCE. And remember that changes can be made to the Red Hat Exam Prep guide at any time.

As of this writing, Red Hat has just added the Network Time Protocol (NTP) to the list of services required for the RHCE. As with other network services, per the RHCE section of the Red Hat Exam Prep guide, RHCEs must be able to do the following:

- Install the packages needed to provide the service.
- Configure SELinux to support the service.
- Configure the service to start when the system is booted.
- Configure the service for basic operation.
- Configure host-based and user-based security for the service.

It also suggests that you need to know how to "diagnose and correct problems with (these) network services" as well as SELinux-related network issues during the Troubleshooting and System Maintenance portion of this exam. As of this writing, the only available SELinux option related to NTP and DHCP is to disable protection for this service.

Some additional protections may be available for individual xinetd services such as rsync and cvs; however, as of this writing, no individual xinetd service is listed in either the Red Hat Exam Prep guide or curriculum for RH133 and RH300.

To establish the connection on a TCP/IP network, a client application needs the IP address of the server and the *port number* associated with the server daemon. All common TCP/IP applications have a standard port number; some examples are shown in Table 13-1.

Table 13-1: Typical Tcp/Ip Port Numbers

Port Number	**Service**
21	FTP
22	SSH
23	Telnet
25	SMTP (outgoing mail)
80	HTTP
443	HTTPS (secure HTTP)
631	Internet Printing Protocol (CUPS configuration)

If you don't specify the port number, TCP/IP assumes that you're using the default port for the specified service. Clients can't connect unless the corresponding server is running on the remote system. If you are managing a server, you may have a number of server daemons to start when Linux is booted.

The **xinetd** (which stands for the Extended Internet Services Daemon) service can start a number of these server daemons simultaneously. The **xinetd** service listens for connection requests for all of the *active* servers with scripts in the /etc/xinetd.d directory. There's a generic configuration file for xinetd services, /etc/xinetd.conf. The scripts in the /etc/xinetd.d directory function as service-specific configuration files.

Generic xinetd Configuration

The generic configuration for xinetd services is stored in the /etc/xinetd.conf file. As RHCE candidates need to configure services only for "basic operation," this chapter analyzes only the active directives in this file. First, a number of default settings are enabled with the following command:

```
d
defaults
```

Certification Objective 13.02-The Secure Shell Package

Red Hat Enterprise Linux installs the Secure Shell (SSH) packages by default, using the openssh-server, openssh-clients, openssh-askpass, and openssh RPMs. The Secure Shell and Secure Copy programs, **ssh** and **scp**, are replacements for the **rsh**, telnet, and **rcp** programs. They encrypt communication between different computers. The secure daemon, **sshd**, listens for all inbound traffic on TCP port 22. The SSH configuration files are located in the /etc/ssh directory.

The Secure Shell daemon works because it encrypts messages. RHEL incorporates SSH version 2, which includes an enhanced key exchange algorithm.

Basic Encrypted Communication

Basic encryption in computer networking normally requires a private key and a public key. You keep the private key and send the public key to others. When they want to send data to you through SSH, their messages are encrypted with the public key. Your computer can descramble the message with the private key.

Encryption keys are based on random numbers. The numbers are so large (typically 512 bits or more) that the chance that someone will break into your system, at least with a PC, is quite small in the foreseeable future. Private and public encryption keys are based on a matched set of these random numbers.

Private Keys

Your private key (essentially a file with your special number) must be secure. When you enable an application, it can attach the key to your messages. Anything you send-say, from your e-mail account-can then be digitally signed and encrypted. The public key is added to the end as part of your signature. Only the recipient will be able to decrypt the message.

Public Keys

Your public key value is just that, publicly available. A central authority such as VeriSign, GlobalSign, or Thawte provides public access to public keys they have created. If they generate a private key for you, they'll keep a secure copy on their system. You can just attach your public key to the e-mail, or the end users can publicly retrieve it from the Web site associated with the central authority.

The example shown in Figure 13-1 lists the directories and files associated with SSH usage as well as a public key that has been added to your "keyring."

```
[root@Enterprise5vm ~]# ls .ssh/
id_dsa  id_dsa.pub  identity  identity.pub  known_hosts
[root@Enterprise5vm ~]# cat .ssh/id_dsa.pub
ssh-dss AAAAB3NzaC1kc3MAAACBAJQ+84NL84wmw3SP61UcE3/Mr8hqYoOFbv3uinYns3EzjWuraYhN
X3bS7zuz7AkGBCQm5PD27GpbEUzSBqDOqPhVraysUXcqAWi/TU0fotFbo6Q4/H8RKLjSy3421G9sRgZk
9YGuk+f5mTw305w8kC5HmoqHhpeT2hTd2Q2DbAs1AAAAFQDny1g13u5FYRgmqEBHi7XfEvnQNwAAAIAM
RhxF9P6pEr4XTKG6SYHM/G17Vm93Fmq7S/q2bVFMRWwy5kP76EI33L2FMB102Um3mvMc76WuLiACpXiq
zkK6Mhxg1HrBwsc/gCrg3R5fUmq3LJDcamn79snf2Q0aZGzAa5S2SNQFps28fsLphIB1Lp1XwagxWE1g
ehqBo5RAYQAAAIA2VU/JYrHR01bOJS46xbkhn1dP3/ZcfcudXA33UVV1rV6Q0xC3beFEQrdQbMxtni4n
jwwXndepqsSxbvssHVjernRduwjrC5RpjF1UPxsGGq/p53xBTchZyQBnAXub0v2K0jiDNgkgrqfd5NGD
4baWLbRxnIhiKXcqMDhayc8z7w== root@Enterprise5vm.example.net
[root@Enterprise5vm ~]# cat .ssh/identity.pub
2048 35 247090728982579742422490392402579479686650097558906089113718203302403196
109415029434527227720641989136324840097385176645535017193319706120786414198440002
440064028844285656188229857410700157886308580560091963698205782347945631228960485
765278319183979458325619920869795411269688877933879295953504488751949751819785556
730177319280542866806876591422849779516923287846174389943921159757224720182964270
823196940365197603112559173382464683050037835088039586619481228333093204590494300
402891277724501204717555659216485762876903368734380440152321774839324870917983740
513624998218528200279604843958790588784496602441643207000762002710 root@Enterpris
e5vm.example.net
[root@Enterprise5vm ~]#
```

Figure 13-1: A public key

This key is like a password used to encrypt your data. Imagine trying to remember the 1024-bit number expressed in hexadecimal value as shown here.

```
3
3 081  89 02 81 81  00D459 6E  01 EB A 01 23CA D51 B7
7
7 835  05A 4 EB FC C7 0B 43 82A733 5 D62A51 B B9 D6
 2
 29EA  860BEC 2B7A B8 2E9 63A 4C 71A 2 D087  11 D0
E
E1 49  4DD5 1E 20 83 82 FA5 8C7 DA  D9 B03 865  FF6E
 8
```

Certification Objective 13.03-Dynamic Host Configuration Protocol (DHCP)

There are two protocols that allow a client computer to get network configuration information from a server: DHCP (Dynamic Host Configuration Protocol) and BOOTP. DHCP works if you have a DHCP server on the local network. The BOOTP protocol is required if you're getting information from a DHCP server on another network.

DHCP servers can simplify and centralize network administration if you're administering more than a few computers on a network. They are especially convenient for networks with a significant number of mobile users. The BOOTP protocol is essentially just a way to access a DHCP server on a remote network.

As of this writing, Red Hat does not include any GUI tool to configure a DHCP server. You'll have to do your work in this section from the command line interface.

Exam Watch

While DHCP knowledge is not explicitly listed in the current Red Hat Exam Prep guide, it is a part of the associated curriculum. Based on their outlines, the RHCT course, RH133, examines the configuration of a DHCP client. The RHCE course, RH300, addresses DHCP servers. It is important for any network administrator to know DHCP. However, it's not in the Red Hat Exam Prep guide; you'll have to make your own decision about whether to learn how to create a DHCP server for your RHCE exam.

Installing DHCP Packages

As with most network services, DHCP has a client and a server. These are based on the dhcp and dhclient RPM packages. The dhclient RPM package should be installed by default; if you're using a service such as NetworkManager, you'll also need the dhcdbd package. If you're working with IPv6, you'll need the dhcpv6_ client. On the server side, the dhcp RPM package is installed by default with the Network Server package group.

On the Job

Red Hat seems to change the commands and packages related to the DHCP client frequently. Older versions of Red Hat have used *dhcpcd* and *pump* as DHCP client commands. Just be aware of this if you're working with an older version of Red Hat.

DHCP Server Configuration

A DHCP server sends messages to multiple computers on a LAN. This is also known as a multicast. It should be enabled by default. You can confirm this with the ifconfig command. The output should resemble Figure 13-4, which includes a **MULTICAST** setting for the active network card.

```
[root@Enterprise5vm ~]# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:0C:29:CE:54:00
          inet addr:192.168.0.22  Bcast:192.168.0.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fece:5400/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1492  Metric:1
          RX packets:13668 errors:0 dropped:0 overruns:0 frame:0
          TX packets:5209 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:12235808 (11.6 MiB)  TX bytes:879604 (858.9 KiB)
          Interrupt:177 Base address:0x1080

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:8716 errors:0 dropped:0 overruns:0 frame:0
          TX packets:8716 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:6047214 (5.7 MiB)  TX bytes:6047214 (5.7 MiB)

[root@Enterprise5vm ~]#
```

Figure 13-4: Active network interfaces *MULTICAST*

If you don't see **MULTICAST** associated with your network card, someone has compiled this feature out of your kernel. For more information on the kernel management process, see Chapter 8.

Now configure the DHCP server daemon, **dhcpd**, by creating or editing the /etc/dhcpd.conf configuration file. Normally, this file allows the DHCP server to assign IP addresses randomly from a specific range. But the default version of this file is blank. You can start with the dhcpd.conf.sample file in the /usr/share/doc/dhcp-*versionnum* directory. The lines that start with a hash mark (#) are comments in the file. Let's analyze this sample file in detail:

- **ddns-update-style interim** With this command, the RHEL DHCP server conforms as closely as possible to the current Dynamic DNS standard, where the DNS database is updated when the DNS server renews its DHCP lease. It is "interim" because the standards for DDNS are not complete as of this writing.
- **ignore client-updates** A good setting if you don't want to allow users on client computers to change their host names.
- **subnet 192.168.0.0 netmask 255.255.255.0** Describes a network with an address of 192.168.0.0 and a subnet mask of 255.255.255.0. This allows the local DHCP server to assign addresses in the range 192.168.0.1 to 192.168.0.254 to different computers on this network. If you've configured a different network IP address, you'll want to change these settings accordingly.
- **option routers** Lists the default router. You can use more than one **option routers** directive if you have more than one connection to an outside network. This information is passed to DHCP clients as the default gateway, which supports access to outside networks such as the Internet. You'll want this command to reflect the IP address for the gateway for your network.
- **option subnet-mask** Specifies the subnet mask for the local network.
- **option nis-domain** Notes the server that provides the NIS shared authorization database. If you've configured NIS on your network, you'll want to substitute the name of your NIS domain for *domain.org*. Otherwise, you should comment out this command.
- **option domain-name** Adds the domain name for your network. Substitute the IP address for the DNS servers you want your clients to use.
-

- **option domain-name-servers** Notes the IP address for the DNS server for your network. You can add more commands of this type to specify additional DNS servers.

- **option time-offset** Lists the difference from Greenwich Mean Time, also known as UTC (a French acronym), in seconds.

- **option ntp-servers** Notes any Network Time Protocol (NTP) servers for keeping the time on the local computer in sync with UTC. I describe NTP later in this chapter.

- **option netbios-name-servers** Adds the location of any Windows Internet Naming Service (WINS) servers for your network. As this is a Microsoft service, I refer to it briefly in the description of Samba in Chapter 10.

- **option netbios-node-type 2** Peer-to-peer node searches, associated with WINS.

- **range dynamic-bootp 192.168.0.128 192.168.0.254** Specifies the assignable IP addresses to *remote* networks, using the BOOTP protocol.

- **default-lease-time** Specifies the lease time for IP address information, in seconds.

- **max-lease-time** Specifies the maximum lease time for IP address information, in seconds.

- **next-server** Notes the boot server for network computers. If you don't have any network computers, you can comment out this entire stanza.

You can also assign a specific IP address to a computer based on a client's Ethernet address. Just add an entry similar to the following to /etc/dhcpd.conf:

```
h
host mommabears {

     hardware ethernet 08:00:12:23:4d:3f;
```

Certification Objective 13.04-The Network Time Protocol (NTP)

When multiple servers are being used for the same purpose, it's important to keep their time clocks in sync. Whether it be a database, timestamped entries for an online business, or backup servers used for load balancing, consistent times are important for smooth operation. This section provides instructions on how to create and configure your own Network Time Protocol (NTP) server for *basic* operation. But first, you'll need to configure a client.

NTP Client Configuration

I'll show you how to configure a client using the GUI. This is one of the few services where client configuration is actually more awkward from the command line. Open the Date/Time Properties tool; one way is to run **system-config-date** from a command line console in the GUI.

I assume the information defined in the Date & Time and Time Zone tabs is already correct. (The information in the tabs should at least be close; the NTP client may have trouble if your clock is off by more than 1000 seconds.) Select the Network Time Protocol tab, as shown in Figure 13-7.

Figure 13-7: Configuring the Network Time Protocol

Make sure the Enable Network Time Protocol option is active. Your administrator (or possibly exam proctor) may tell you to use a different NTP server, which can be changed in the NTP Servers text box. If you want to find a public NTP server other than the defaults shown, refer to http://ntp.isc.org/bin/view/Servers/WebHome. Generally, you should only connect to a "Stratum Two" time server. Even then, you'll generally need to ask permission of the NTP server administrator before connecting your time server. Too many connections to an NTP server can degrade performance, leading to delays. And delays are never good for a time server.

On the Job

One option to Stratum Two servers is available from the Public NTP Time Server project, available at www.pool.ntp.org. Red Hat's default time servers for RHEL 5 are part of this project.

Once you configure a local NTP time server, you can set other NTP clients to synchronize with that server.

Under the Advanced options, unless there are problems with the initial network connection, you should select Synchronize System Clock Before Starting Service. On the other hand, this option can slow the boot process, which can be a problem during an exam.

If you select Use Local Time Source, it adds the following directive to /etc/ntp.conf:

```
s
server 127.127.1.0
```

Certification Summary

Networking services are an integral part of Red Hat Enterprise Linux. While most network services have been covered in other chapters, a few just don't fit neatly into other categories. Some are controlled through the xinetd daemon. SSH and DHCP are two other network services.

The Extended Internet Services Daemon, xinetd, governs the services configured though the /etc/xinetd.d directory. It controls services such as rsync, gssftp, and krb5-telnet, when installed.

Any network that is connected to an insecure network such as the Internet is vulnerable. Older protocols such as Telnet allowed passwords in clear text, which can be easily captured with tools such as Wireshark. The OpenSSH server can help you set up encrypted communication between computers.

DHCP allows a network administrator to manage IP address assignments of the computers on a LAN from a centralized server. DHCP requires some specialized setup on both the client and the server; however, it is easy to maintain once it is configured.

NTP services support time synchronization, which can keep changes such as moving to backup servers, load balancing, logs, and more, seamless to the end user. Clients can be configured with the Date/Time configuration tool or in /etc/ntp.conf. Configuring an NTP server requires specialized settings in /etc/ntp.conf.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 13.

The Extended Internet Services Daemon (xinetd)

? xinetd acts as a "super-server" for a number of other network services, such as the Kerberos secured versions of Telnet and rsync.

? Individual services have their own management scripts in the /etc/xinetd.d directory.

? Most xinetd services are disabled by default.

? You can activate an xinetd service with the appropriate chkconfig command or by directly editing its xinetd script.

? xinetd listens for connection requests from client applications.

? When xinetd receives a connection request, it starts the server associated with the TCP/IP port and then waits for other connection requests.

The Secure Shell Package

? The OpenSSH command utilities-**sshd**, **ssh**, **ssh-keygen**, **ssh-add**, and **ssh-agent**-provide secure remote services over any network connections.

? Encryption is based on private and public keys.

? Public keys are shared with others, so that they can communicate with you through SSH.

? As it is easy to decipher traffic, even passwords, from telnet, **ftp**, and the "r" commands, it is best to use SSH on any publicly accessible network.

Dynamic Host Configuration Protocol (DHCP)

? DHCP allows a client computer to obtain network information (such as an IP number) from a server.

? The BOOTP protocol allows a client computer to access a DHCP server on a remote network.

? DHCP servers are configured through /etc/dhcpd.conf.

? The DHCP server daemon is **dhcpd**; the DHCP client daemon is **dhclient**.

The Network Time Protocol (NTP)

? NTP servers can help synchronize the systems on a network.

? NTP clients can be configured with the Date/Time Configuration tool, which can be started with the **system-config-date** command. Alternatively, clients can also be configured in /etc/ntp.conf.

? NTP clients can be synchronized with Stratum Two servers with the permission of their administrators; one alternative is the Public Time Server project at pool.ntp.org.

? To configure an NTP client as a server, the /etc/ntp.conf file needs to be configured to allow access to the desired networks.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer for many of these questions.

The Extended Internet Services Daemon (xinetd)

1. You are using the xinetd program to start services. What command makes sure the Kerberos version of Telnet, as configured in /etc/xinetd.d/krb5-telnet, starts the next time you boot Linux? Assume the xinetd service is active. ?

   ____

2. What command rereads any configuration changes you've made to a file in the /etc/xinetd.d directory? ?

   ____

Answers

1. The command that makes sure that the Kerberos version of Telnet starts the next time you boot Linux is **chkconfig krb5-telnet on**. This assumes the xinetd service is active when you boot. Naturally, you can edit /etc/xinetd.d/krb5-telnet directly for the same result.

2. The command that rereads any configuration changes you've made to a file in the /etc/xinetd.d directory is **service xinetd reload**.

The Secure Shell Package

3. Once configured, what do you have to do to an existing firewall to allow Secure Shell (SSH) access? ?

   ____

4. You've started the SSH service and have a local account under username troosevelt. The local computer name is conservation. How would you log into that account from a remote computer? ?

   ____

5. Where are SSH keys stored for user michael? Assume a standard home directory. ?

   ____

6. When configuring an SSH server, what directive prohibits SSH logins by the root user? ?

   ____

7. When configuring an SSH server, what directive supports access to remote GUI applications?

____

8. If you've configured access to remote GUI applications, how would you log into the account described in question 4?

____

Answers

3. Once configured, you need to open TCP port 22 to allow Secure Shell (SSH) access. It's easy to do with the Red Hat Security Level configuration tool described in Chapter 15.

4. To log into the troosevelt account from a remote computer, use the **ssh troosevelt@conservation** command.

5. SSH keys for user michael are stored in the /home/michael/.ssh/directory.

6. When configuring an SSH server, the directive that prohibits SSH logins by the root user is **PermitRootLogin no**.

7. When configuring an SSH server, the directive that supports access to remote GUI applications is **X11Forwarding yes**.

8. If you've configured access to remote GUI applications, you would use the **ssh -X troosevelt@conservation** command to log into the account described in question 5. You can then access GUI applications such as Firefox and OpenOffice.org writer through the SSH connection.

Dynamic Host Configuration Protocol (DHCP)

9. What would you add to the /etc/dhcpd.conf configuration file if the DNS server you want your clients to use has an IP address of 10.11.12.1?

____

10. You add a new workstation to your /etc/dhcpd.conf file. You're in a hurry to finish, so you save and go to lunch. When you return, your phone mail is full of user complaints that they can't access the Internet, but the local network is fine. You surmise that you accidentally changed something in the dhcpd.conf file. What directive should you look at first in dhcpd.conf?

____

Answers

9. If the DNS server you want your clients to use has an IP address of 10.11.12.1, add the **domain-name-servers 10.11.12.1;** directive to your /etc/dhcpd.conf configuration file.

10. A missing **routers** directive in /etc/dhcpd.conf would keep your hosts from getting the gateway address, which is required to access the Internet from an internal network.

The Network Time Protocol

11. From the /etc/ntp.conf file, there's the following directive:

```
restrict default kod nomodify notrap nopeer noquery
```

12. What Web site is associated with the public NTP server project?

Answers

11. The appropriate directive that limits access to the 192.168.0.0/24 network is

```
restrict 192.168.0.0 mask 255.255.255.0 nomodify notrap
```

12. Pool.ntp.org is the Web site associated with the Public Time Server project. If you've used the Date/Time Properties tool, you should be familiar with this from the Red Hat defaults.

Lab Questions

Lab 1

1. Your company bought another competitor on the opposite coast recently, just as the new corporate application was being deployed everywhere, so you sent the app to them, too. They use a Unix host for this application on their network. You need to be able to connect to this host *securely* for maintenance purposes on the new system-wide application you deployed. Both networks have Internet access.

   What do you do?

Answers

1. If you need access now, and both systems are connected to the Internet, you can set up SSH for secure communications. If the other network does not already have it installed, have the administrator in the remote location download and install it, and then create an account for you.

   The basic steps outlined here may vary with the version of Unix used on the other network. If it's Red Hat Enterprise Linux, all you need are the SSH packages described in this chapter, but you may not have that luxury.

   Download the source code for OpenSSH and put it into a specific directory. Assuming your version of Unix can't handle RPMs, you'll need to unpack a tarball. You can then unpackage the files in the tarball and use the files in the resulting directory to compile and configure a Secure Shell server. Once it is configured, you have the option to set up private and public keys.

   If you don't need immediate access, you could, alternatively, configure a computer with Red Hat Enterprise Linux and a Secure Shell server. Send the computer to the administrator of the remote Unix network. Have the admin add it to his or her network, and you can check the problem from your site securely. The application is running on the Linux computer that you sent. (Alternatively, you can set up OpenSSH on Microsoft Windows, as described earlier in this chapter.)

Lab 2

2. You'll also need two Linux computers for this lab: one as a DHCP server and a second as a DHCP client. Using the DHCP server created earlier in this chapter, set up a static IP address for the computer of your choice. You'll want to assign a specific name for that server, precious.example.com, and a special IP address on the 10.11.12.0 network, 10.11.12.13. Assume that you've already set up the example.com network as well as an appropriately configured DNS server.

Answers

2. Assuming you've read the chapter, you've seen the template in the dhcpd.conf.sample configuration file for a static IP address:

```
host ns {
    next-server marvin.redhat.com;
    hardware ethernet 12:34:56:78:AB:CD;
```

Lab 3

3. To do this lab, you'll need the help of a partner. Have him or her set up your system as described in the answer to this lab in the next section. The intent of this lab is to help improve your troubleshooting skills for the Troubleshooting and System Maintenance section of the RHCE exam. ?

Answers

3.

Lab 3: Part 1

You're going to set up this lab for your partner (I've set up these "answers" in a different order, first the setup of the lab, to help discourage "shoulder surfing"), using the following steps:

1. Make sure your system supports the SSH service, is active, and includes an account for your partner's username and password. If the password is secret, let your partner enter the password.
2. Take over your partner's RHEL system. Make sure there's a connection between your computers.
3. Configure the SSH service as described in this chapter. If you have a firewall configured on this system, make sure to open port 22 to allow communication to the local SSH service. You can use the Red Hat Security Level configuration tool described in Chapter 15 to help.
4. Modify the SSH server configuration file to allow users, not including the regular username for your partner. As described in this chapter, this involves the **AllowUsers** directive in the /etc/ssh/sshd_config file.
5. Don't forget to make the sshd service reread the configuration file with a command such as **service sshd reload**.
6. Pass your partner's system back to him or her. Instructions for your partner can be found in Part 2 of Lab 3.

Lab 3: Part 2

1. Take your system back from your partner. Log in as a regular user. Use the **ssh** command to log into your partner's system under your username. Repeat the process to log back into your own system. What happens?
2. Return to your own system. Analyze the logs. What do you see? Is there anything special in /var/log/secure?

Lab 4

4. This lab includes a matching scenario to Lab 3 for your partner. Refer to the "Lab Answers" section for what you need to configure. ?

Answers

4.

Lab 4: Part 1

1. Make sure your system supports the SSH service, is active, and includes an account for your partner's username and password. If the password is secret, let your partner enter the password.
2. Take over your partner's RHEL system. Make sure there's a connection between your computers.
3. Configure the SSH service as described in this chapter. If you have a firewall configured on this system, make sure to open port 22 to allow communication to the local SSH service. You can use the Red Hat Security Level configuration tool described in Chapter 15 to help.
4. Modify the SSH server configuration file to prohibit logins by the root user. As described in this chapter, this involves the **PermitRootLogin no** directive in the /etc/ssh/sshd_config file.
5. Don't forget to make the sshd service reread the configuration file with a command such as **service sshd reload**.
6. Pass your partner's system back to him or her. Instructions for your partner can be found in Part 2 of Lab 4.

Lab 4: Part 2

1. Take your system back from your partner. Log in as a regular user. Use the **ssh** command to log into your partner's system as the root user. Repeat the process to log back into your own system. What happens?
2. Return to your own system. Analyze the logs. What do you see? Is there anything special in /var/log/secure?

Lab 5

5. This lab requires two Red Hat-capable systems, or at least a second system that you know how to configure as an NTP client. Configure the RHEL system as an NTP server. Allow access on your local IP address subnet. Make sure the service is started and will start the next time you boot Linux.

Answers

5. Before configuring a system as an NTP server, you have to first configure it as a client. The simplest method for doing so is to use the Date/Time Configuration tool, which can be started in the GUI with the **system-config-date** command.

   Once configured as a client, make sure the NTP service is running and set to start the next time Linux is booted on the local system. The simplest method involves the following commands, which should be familiar if you've configured other services on Red Hat systems:

```
#
# chkconfig ntpd on
#
```

PREV
NEXT

Chapter 14: The X Window System

Overview

One of the most important aspects of getting a Linux system up and running is configuring the user interface. As RHCEs and RHCTs are expected to configure computers for non-administrative users, the Red Hat exams test your ability to configure the X Window System, which is the foundation of the Linux graphical user interface (GUI). While the GUI plays an integral part of other operating systems such as Microsoft Windows, the X Window System on Linux is essentially just another application.

Many administrators don't even bother with the GUI; the command line interface is enough for most administrative purposes. However, regular users on a Linux workstation are more productive using the GUI and the multitude of X Window-based applications. If you are helping users migrate from Microsoft Windows to Linux, the X Window System allows you to provide a less intimidating environment.

Not all Linux computers require the X Window System. For example, computers that are used as dedicated DHCP, DNS, or NFS servers generally don't serve as workstations for anyone and therefore don't need any sort of GUI. Many Linux gurus are biased against the GUI. While Red Hat and others have developed some helpful GUI tools, they are almost always "front ends," or programs that customize one or more commands at the command line interface.

But if you're administering a network of Linux computers for regular users, you'll need to know how to administer the X Window System, a skill that requires a basic understanding of the available desktops and graphical applications.

Most Linux distributions (including RHEL) have converted display software from the XFree86 to the X.org system. While the names of the configuration files and some of the commands have changed, the basic settings and tools have not. If you learned to configure the X Window using XFree86, you should have no trouble configuring the X.org system.

This chapter starts with the X server, as configured on the local computer. It continues with X clients, as generic applications that you can run from the local or a remote network computer. Once everything is configured, you're ready to take a step back to the start process for the X Window. The chapter moves on to the two major Linux graphical desktops, and finally covers a very few of the available graphical applications.

Inside the Exam

The Linux Graphical User Interface

The Red Hat Exam Prep guide suggests that you need to know how to configure the X Window, presumably for non-administrative users. Remember that RHCE candidates must successfully complete *all* RHCT Troubleshooting and System Maintenance requirements, including configuring the X Window System and a desktop environment. You also need to know how to configure the X Window on a local computer. There are a number of reasons why the X Window may fail.

The X configuration files can be difficult to learn. It may be more efficient to use the Red Hat GUI X Window Display Settings configuration tool, which you can start with the **system-config-display** command. While Linux geeks generally shy away from GUI tools, you need to use the system that works most quickly for you.

The X Window System can work over a network. Once properly configured, you can run GUI applications from a remote computer. To make this work, you need to understand modularity of the X server and X clients, as well as the

way X Window security is managed on your network.

Exam Watch

I use the Red Hat Display Settings tool and *system-config-display* command interchangeably; the command is the fastest way to start the tool.

Certification Objective 14.01-X with Clients and Servers

The X Window System is designed as a flexible and powerful client/server-based system. To configure and troubleshoot the X Window interface, you need to understand the client/server nature of the X Window System.

As you might have guessed from the terms client and server, the X Window System is designed to work in a networked environment. The client and server can both reside on your own computer or on separate computers on the network. In other words, not only can you run X applications on your system, you can run X applications on other computers on your network. The graphical displays from those remote applications are sent to your monitor.

In fact, X Window applications handle this task so well that, providing the network is fast enough, you really can't tell from a performance point of view which applications are running locally and which applications are running remotely.

When you configure the X server, I'll show you the modularity of the system. In brief, components such as keyboards, mice, and monitors are configured separately and all become modular components of the X server. While one X server process controls the display, you can run as many X clients as your hardware resources, primarily RAM, will support. If your Linux system is part of a network, you can also start X clients on other systems on the network and have those clients send their displays to your X server.

X clients exist for almost every basic application-word processing, spreadsheets, games, and more. The Red Hat GUI configuration tools were developed as X clients. There are even X client versions of popular utilities such as the emacs editor.

Different Meanings for Client and Server

Normally on a network, the local computer is the client and the remote computer acts as the server. X Window clients and servers work on a different paradigm. The X server controls the graphics on the local computer. The X server draws images on your screen and takes input from *your* keyboard and mouse. In contrast, X clients are local or remote applications such as **xclock** that you can run on the local X server.

You can run an X client locally or remotely. Local X clients run on your workstation; remote X clients run on the local X server. When you run a remote X Window client application, you start the program on a different computer and send its output to use the X server on your local computer. Figure 14-1 illustrates a local X server with one local and one remote X client.

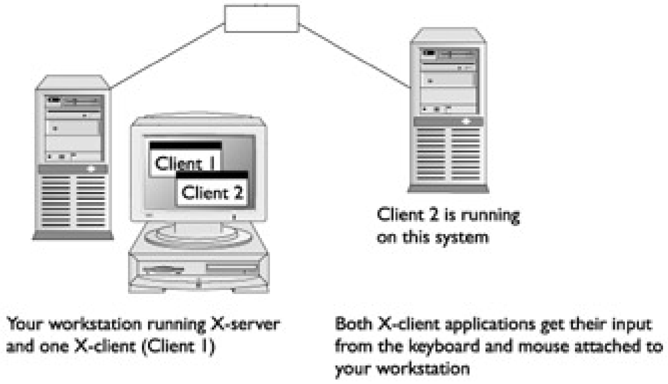

Figure 14-1: Running X Window clients from remote or local computers

Before reading more about running X client applications, you need to first configure an X server. While RHEL normally configures most hardware configurations during the installation process, you need to know how to edit the X

Window configuration file, directly in a text editor or with the Red Hat Display Settings tool (**system-config-display**).

On the Job

If you're in the GNOME (GNU Network Object Model Environment) or KDE (K Desktop Environment) desktops, you can also start the Red Hat Display Settings tool with the GNOME System (or KDE Main Menu) | Administration | Display command.

Supported Hardware

If a Linux installation program does not successfully configure the X Window System, it can be difficult to configure the GUI on a Linux workstation. Fortunately, RHEL comes with tools and drivers that make this job relatively painless and easy.

But it all depends on the hardware. Linux may not always work with the latest and greatest video card or monitor. Many video cards and monitors include proprietary software; it may take some time before Linux developers are able to "reverse-engineer" these components. For the latest official information, check the Linux Video and DVD Project at www.linuxvideo.org, the Linux Hardware Compatibility List at www.tldp.org/HOWTO/Hardware-HOWTO/, and the Red Hat Hardware Catalog at https://hardware.redhat.com/hwcert/.

On the job

Linux provides world-class support for graphics. The list of movie studios that use Linux to create feature films is impressive (including Disney, DreamWorks, Industrial Light and Magic, Paramount, and so on). If you need motion picture-quality graphics support, you may want to consider commercial alternatives to the X.org server, such as Accelerated-X from Xi Graphics (www.xinside.com).

The X Window server shipped with RHEL 5 is an open-source X server program from X.org. This server supports hundreds of video cards and monitors. The best places to check to determine whether your video card and monitor are supported are the hardware lists described earlier. Alternatively, navigate to the X.org Web site to find the latest support information.

The latest version of X.org includes modules for different video servers. Hardware support for most video servers is already there. If you learn of updates, changes are easy. Just add the module, and then point to it in the /etc/X11/xorg.conf configuration file.

If you are using an unsupported video card, support is also included for generic VGA and VESA devices. Most video cards and monitors will work with these X servers.

Default X Clients

When you configure a workstation for most users, you'll be configuring their GUI. You may need to specify a default desktop such as GNOME or KDE. I'll describe these desktops later in this chapter. You may want to set up specific icons for that user's desktop. Settings for default X clients are stored in each user's home directory, in various hidden directories. In RHEL, KDE configuration files are stored in each user's ~/.kde directory (where the tilde, ~, represents a Linux home directory). GNOME configuration files are stored in hidden directories such as ~/.gconf, ~/.gnome, and ~/.gnome2.

Exercise 14-1: Starting X Server

In this exercise, you will start your X server without a window manager. You'll then start an X client application

known as **xterm**. Some of the commands used in this exercise are covered later in the chapter. If the X Window System is not running, you can skip steps 1 and 3.

1. If the X Window System is running, change to a text console by pressing CTRL-ALT-F1.
2. If you see a login prompt, log in at the text console as the root user. Otherwise, press CTRL-C to stop the X Window.
3. When you log in as root at the text console, stop the current X Window server with the following command (**telinit 3** works as well):

Certification Objective 14.02-The X.org Server Configuration

Most configuration files associated with the X.org server can be found in the /etc/X11 directory. While the focus is on the xorg.conf file, there are other files in that directory. The way the X.org server is designed, you can start the X Window in multiple terminals. You can boot directly into the X Window in your choice of display managers or use **startx** to start it from a command line console.

X.org Server Configuration Files

A wide variety of X Window configuration files are located in the /etc/X11 directory. Many are discussed in other parts of this chapter. While I don't cover these files in detail, the Red Hat Exam Prep guide doesn't say much about configuring the X Window in detail.

Therefore, I focus here on the primary X Window configuration file, xorg.conf. It's instructive to read the associated man page carefully. It is well documented and includes a number of commented sample commands that can help you configure your system in a number of special ways. For example, it includes tips on how you can

- Configure different keyboards.
- Set up multiple monitors, in what is known as a "multi-head" configuration.
- Disable switching from the GUI to other virtual terminals with the following command:

Certification Objective 14.03-Tools for X.org Configuration

If you want to configure your X Window System, there are three options: Direct configuration of the X Window configuration file, the Red Hat Display Settings tool (**system-config-display**), or automatic installation and configuration of the X Window during the installation process.

Even if you didn't install any graphics software when you installed RHEL, you can still use the **system-config-display** command. It starts its own default graphics mode if it detects a graphics driver.

Red Hat Display Settings Tool

The Red Hat Display Settings tool is a stand-alone program that you can run at any time from the command line. The basic routines that start with the **system-config-display** command are also used by the Red Hat installation program if you choose to install and configure the X Window System at that time.

The **system-config-display** program is a character-based menu-driven interface that helps you configure your video hardware. If you're starting from a text console, it automatically probes your video card and selects the appropriate X server image. If **system-config-display** cannot detect your graphics card, it allows you to select it from the list of supported video cards.

It's easy to start the Red Hat Display Settings tool. Just type **system-config-display** at a command line interface. It provides a simple GUI, even if you start it from a regular text console. It starts the Display Settings window similar to that shown in Figure 14-7.

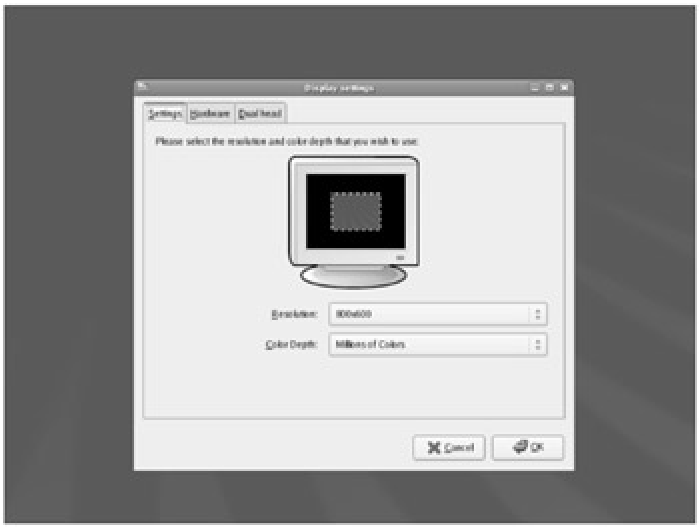

Figure 14-7: The Display Settings tool, started from the text console

You can set the default resolution and color depth under the Settings tab. If the Display Settings tool successfully identifies your hardware, you'll see it listed under the Hardware tab. In this case, it detected a VMware graphics driver with a LCD monitor. You can change these settings by clicking the associated button. If your hardware supports it, you can configure:

- Monitor resolutions between 640×480 and 1920×1440.
- A color depth of thousands or millions of colors. Thousands corresponds to 16-bit color, and millions

corresponds to 24- or 32-bit color, depending on the capability of your hardware.

But if you want to select a different hardware component, you can select it from a list. Click the Hardware tab, shown in Figure 14-8.

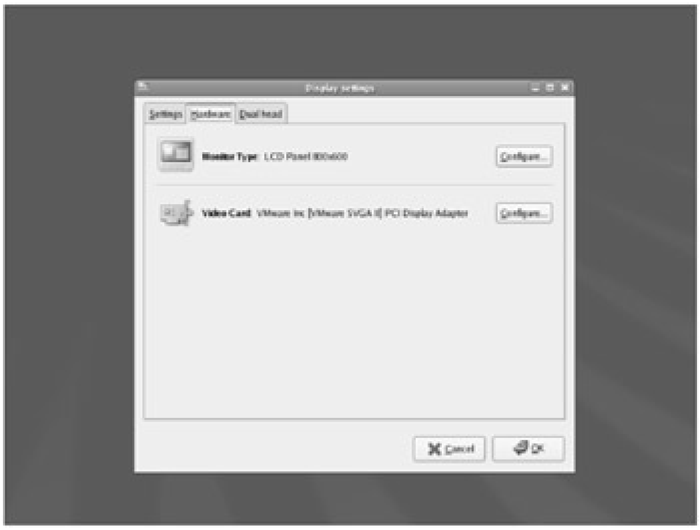

Figure 14-8: Display settings

You'll see options to configure your monitor and video card. Click the Configure button in the Video Card section. This should bring up the Video Card dialog box shown in Figure 14-9.

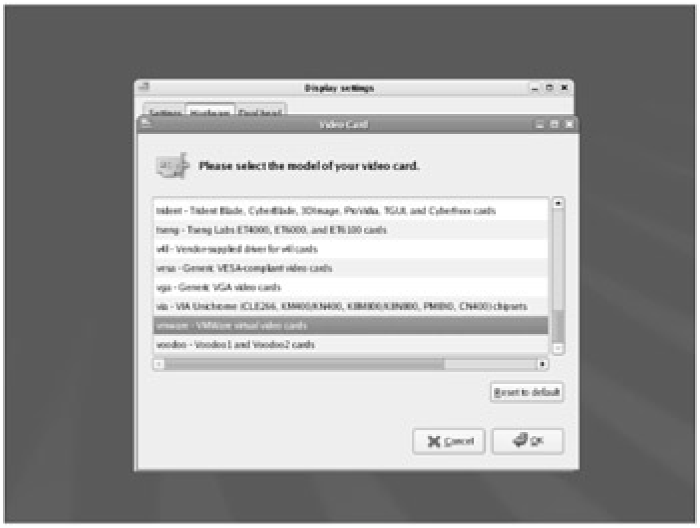

Figure 14-9: Selecting a graphics card

Browse through the list of video cards. If you do not see your graphics card here, it may not be supported. In this case, there are several options:

- Select a video card at least similar to your model. Alternatively, you may find a generic server such as a VESA driver (generic) that is compatible with your video card. Test and if necessary edit the /etc/X11/xorg.conf file to complete your changes.
- Check online for other Linux users who are running the X Window System with the same type of hardware. Red Hat maintains mailing lists where many users ask such questions; for the current set of lists, see www.redhat.com/mailman/listinfo.
- Use the Unsupported VGA or VESA compatible X Window Server.

On the Job

The VESA (Video Electronics Standards Association) driver is also known as SVGA (Super Video Graphics Array).

- Go to www.X.org and download the latest drivers. You'll need to edit the /etc/X11/xorg.conf file directly to point to this driver.

Once your selections are complete, click OK. This returns you to the Hardware tab. If your video card is so capable, you'll be able to activate the Enable Hardware 3D Acceleration option. (Otherwise, you won't even see the option in the menu.) Next, configure the Monitor. Click the Configure button in the Monitor Type section of the Hardware tab to open the dialog box shown in Figure 14-10.

Figure 14-10: Selecting a monitor

On the Job

Configuring the X Window System to run on a laptop can create challenges. If you are planning to install Red Hat Enterprise Linux on a laptop, a good source for tips and additional information is the Linux on Laptops Web site at www.linux-laptop.net. When all else failed during a recent installation, I was able to run *system-config-display* remotely, over an SSH connection.

Other Available Tools

There are a number of ways to configure the X Window with command line tools. If you have problems with the X Window, the first thing to do is check the log file, in /var/log/Xorg.0.log. More information is available in the "Troubleshooting" section later in this chapter.

For a quicker look at the start process, move to runlevel 3 with the **init 3** command. Simulate the GUI start process with the following command:

```
#
# Xorg -probeonly
```

Certification Objective 14.04-Running Remote X Applications

One of the most powerful features of the X Window System is its networking support. The X Window System was designed to run in a networked environment. If you are a system administrator responsible for a number of RHEL systems, you don't have to run to the server room every time you want to run a GUI administration tool. With the X Window System, you can connect to any number of systems and redirect the output from X clients running on those systems back to the X server running on your desktop.

One method for running remote X applications was described in Chapter 13. It required appropriate configuration of SSH on the remote X client system, permission on the firewall, and appropriate login (with the **ssh -X** or **ssh -Y** command, followed by ***login@remotepc***).

Exercise 14-4: Starting a Display from a Remote Client

[redacted] exercise, examine the steps required to run a display from a remote client. This assumes a basic knowledge of the Secure Shell and its associated command, **ssh**. For more information on running the **ssh** command, see Chapter 13. You'll need two computers running Linux. While RHEL is not required on both computers, you will need the Secure Shell installed on both. The Secure Shell service, **sshd**, should be running on the remote computer. You'll also need the root password for both computers. In my example, the local computer is named Enterprise, and the remote computer is named cosmicc (which is short for Cosmic Charlie, my favorite Grateful Dead teddy bear). Substitute the names (or IP addresses) of your computers accordingly.

1. On the Enterprise computer, start the X Window. If it isn't already open, use the **startx** command.
2. When the Linux GUI is open, access a new terminal. In GNOME, click Applications | Accessories | Terminal. (In KDE, click Main Menu | System | Terminal.)
3. Log into the remote computer using the Secure Shell. To log in as root, use the following command. Enter the root password on the remote computer when prompted. If you're asked if you want to set up a encryption key, type **yes**. This should log you into the remote computer.

Certification Objective 14.05-Desktops and Window Managers

Part of the Linux GUI is a special type of X client known as a window manager. Earlier in this chapter in Exercise 14-1, you started the X.org server with the **X** command (which is linked to /usr/bin/Xorg). It turned your display into a blank electronic canvas. This is the default desktop display for the X.org server, which is an uninteresting textured gray background. The default mouse pointer for the X Window display is a graphic representation of an *X*.

Once the X.org server starts and this canvas is on your screen, the X server is ready to start serving X clients.

Still, you don't have any of the useful features that you've come to expect in a GUI, such as borders, title bars, menu bars, and minimize/maximize buttons. For this purpose, you need a window manager. A window manager is a special type of X client that can run only with an X server. The window manager controls how other X clients appear on your display. This includes everything from placing title bars and drawing borders around the window for each X client application you start, to determining the size of your desktop. In a nutshell, the window manager controls the look and feel of your GUI.

As is usually the case with all things Linux, you have multiple ways to do the same thing. RHEL can be installed with several different window managers and desktops. The GNOME and KDE desktops include their own window managers. Your choice of window manager and desktop will drive the look, feel, and functionality of the Linux X Window System.

The GNOME and KDE Desktops

Two powerful virtual desktop environments that come with RHEL are the GNOME Desktop Environment and KDE Desktop Environment. The GNOME desktop, shown in Figure 14-11, is the default desktop for RHEL that you first see after installing the X Window System. The KDE desktop, shown in Figure 14-12, is the main alternate desktop system. KDE is the default for several other Linux distributions.

Figure 14-11: The GNOME desktop

Figure 14-12: The KDE desktop

GNOME Features

The GNOME desktop includes support for the GTK+ (GIMP) toolkit, which allows GNOME software components written in any language and running on different systems to work together. In addition, GNOME includes support from a number of other projects, including GConf, ORBIT, and more.

Using GNOME

Many of the features of the GNOME interface will be familiar to you from other desktop environments. On the left side of the screen are icons representing files and applications that you can be open by double-clicking them with the mouse. The GNOME Desktop Environment also provides several virtual desktops. Next to the application buttons on the right side of the panel is a pager you can use to move from one area of the desktop to another.

One of the key features of GNOME are the *panels,* which you can see at the top and bottom of the screen in Figure 14-11. These panels are the control centers for most of your activities while you use GNOME. The button at the far left of the top panel with the imprint of a red hat is the Applications button. Click this button, and you will see a list of menus and submenus that start applications. The button associated with the System menu launches a similar submenu; The System | Administration submenu opens a number of interesting administrative applications.

GNOME includes a number of applications, including graphics tools and an office suite, GNOME Office. As the default Red Hat desktop is GNOME, the remainder of this book will be based on this desktop environment. Nevertheless, the Red Hat exam requirements do not specify a preferred desktop; you should have no problems using KDE or the command line console to do everything that is required for the exam. You may be asked to configure either desktop on the Red Hat exams.

If you configure the default GNOME desktop for your users, you may want to configure GNOME in a special way. Normally, GNOME opens with a number of icons and possibly default applications such as nautilus. You can add more default applications such as a new terminal window or applets such as the **xminicom** modem manager with the Sessions tool, which you can access via the System | Preferences | More Preferences | Sessions command.

KDE Features

The KDE desktop is built on the Qt C++ cross-platform GUI toolkit. This is another versatile way to create GUI applications for Linux.

Many of the features of KDE should also be familiar to you from other desktop environments. In fact, you can configure KDE to a look and feel that is quite similar to Windows 9*x*/2000/XP/Vista. As shown in Figure 14-12, it includes a Main Menu button, represented by the Red Hat in the lower-left corner of the desktop. Like GNOME, it

can include pagers and buttons representing the open programs on the desktop. However, the default version of the KDE desktop is pretty empty, which does not illustrate the capabilities of KDE, and I'm guessing displeases most KDE loyalists.

Default Desktop

Once you've configured the X Window, it's easy to start a Linux GUI. If it isn't already configured to start automatically, run the **startx** command. This command, in the /usr/bin directory, calls configuration files from your home directory. If these files don't exist, they are taken from the default directory for GUI configuration, /etc/X11.

To manage the default desktop, use the **switchdesk** command. It's also known as the Desktop Switching Tool; as they're not installed by default, you may need to first install at least the switchdesk RPM-and if you want to use the GUI version, the switchdesk-gui RPM packages. For example, the following commands set the default desktop to KDE and GNOME, respectively:

```
#
# switchdesk KDE
#
```

Certification Summary

The X Window System provides a state-of-the-art graphical user interface and offers features not found in other GUI environments. Although the X Window System can be complicated, you should be able to configure it during the RHEL installation process. Alternatively, the Red Hat Display Settings tool simplifies the setup or reconfiguration process.

One of the key parts of the X Window System is the X Font Server. If it isn't running, or the /tmp or /home directory partitions are full, you can't run the Linux GUI.

The X Window System works as client and server. X clients and X servers can be located on different computers on a network. With an appropriate firewall configuration, you can take advantage of it with the Secure Shell.

The look and feel of the X Window interface is determined by your choice of desktop. RHEL comes with several desktop environments, including GNOME, KDE, and twm; the default is GNOME.

You can customize the GUI start process in a number of ways. You can configure a default login manager. With an appropriate .xinitrc file, you can set up X clients to start when you run the **startx** command. Plus, you can configure a default desktop with the **switchdesk** command, which defined it in the local .Xclients-default file. Alternatively, you can run the Sessions utility to configure X clients in the GNOME desktop.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 14.

X with Clients and Servers

- ? The X server software manages the graphics display on the local computer, which includes your monitor/graphics adapter, keyboard, and mouse.
- ? You can configure the X server during the installation process. You can also configure or modify the configuration using the Red Hat Display Settings tool.
- ? The X server configuration is stored in the /etc/X11/xorg.conf file.

The X.org Server Configuration

- ? There are a variety of X.org server configuration files in the /etc/X11 directory.
- ? You can start X servers in multiple consoles.
- ? You can configure the X Window to boot directly into a GUI login manager, or start it from a text console.
- ? When you start the X Window, you can customize the process.

Tools for X.org Configuration

- ? The main X.org Server Configuration utility is the Red Hat Display Settings tool, which you can start in either the GUI or a text console with the **system-config-display** command.
- ? Other X.org command line tools are available, such as **Xorg -configure** and **Xorg -probeonly**.

Running Remote X Applications

- ? By default, X clients sends display output to the local computer.
- ? You can log into the remote client using the Secure Shell; if you log in with the **-X** or **-Y** switch, you can send remote X clients back to your local system.

Desktops and Window Managers

?	The X Window System gives you a blank electronic canvas. The look and feel of a GUI is provided by the window manager and desktop.
?	The two main desktop environments are GNOME and KDE.
?	You can use **switchdesk** from a terminal window or the command line interface console to select your default desktop.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

X with Clients and Servers

1. If you're running a GUI tool from a remote system, where is the X client?

2. What is the main X configuration file? In what directory?

3. What command starts a GUI with a blank screen?

Answers

1. If you're running a GUI tool from a remote system, the X client comes from the remote system. The X server is local and manages the hardware.

2. The main X.org configuration file is xorg.conf, in the /etc/X11 directory.

3. The main command that starts the GUI with a blank screen is **X**; as it's linked to the Xorg command, you can use it as well. Yes, with appropriate changes to configuration files, you can start the GUI in a blank screen with other commands, most easily in twm.

The X.org Server Configuration

4. What directive in xorg.conf is associated with keyboards and mice?

5. What command is a text script that starts the GUI from a command line console?

6. What configuration file in /etc/X11 can you use to specify a default GUI login manager?

Answers

4. The directive in xorg.conf associated with keyboards and mice is **InputDevice**.

5. The text script command that starts the GUI from a command line console is **startx**. You can edit it directly in the /usr/bin/ directory.

6. The configuration file in the /etc/X11 directory that you can use to specify a default GUI login manager is prefdm; the key directive in this file is **preferred**.

Tools for X.org Configuration

7. What command tests the X.org start sequence, without actually starting the X Window?

____

8. If you've logged into the command line console in runlevel 3, what command allows you to start the graphical X configuration tool?

____

9. If you've logged into the command line console in runlevel 3 and can't start the graphical X configuration tool, what command creates a local xorg.conf file?

____

Answers

7. The command that tests the X.org start sequence, without actually starting the X Window, is **Xorg -probeonly**. (As you can confirm in the Xorg man page, only one dash is required for Xorg command switches.)

8. If you've logged into the command line console in runlevel 3, the command that allows you to start the graphical X configuration tool is **system-config-display**.

9. If you've logged into the command line console in runlevel 3 and can't start the graphical X configuration tool, you can create a local xorg.conf file with the **Xorg -configure** command.

Running Remote X Applications

10. You've configured the Secure Shell on both systems. If you're logging into the cosmicc.example .net system as user michael, what command would you use to gain access to X applications on the remote system?

____

Answers

10. You've configured the Secure Shell on both systems. To log into the cosmicc.example.net system as user michael, with access to GUI applications, run either the **ssh -X michael@cosmicc.example.net** or **ssh -Y michael@cosmicc.example.net** command.

Desktops and Window Managers

11. What command switches the default desktop environment to KDE?

____

12. In the future, you want to boot Linux directly into the GUI. How would you make this happen in /etc/inittab?

____

Answers

11. The command that sets the default desktop to KDE is **switchdesk kde**.

12. In the future, to boot Linux directly into the GUI, make sure /etc/inittab includes the following directive, which boots Linux into runlevel 5:

```
i
id:5:initdefault
```

Lab Questions

Lab 1

1. You want to upgrade the video card in your Linux system. Your old video card is slow and doesn't have enough display memory to provide you with the resolution and color depth you require. You have obtained a new ATI 32MB Radeon card (I'm using this product for example purposes only). What steps might you follow to replace your old card with your new card?

Answers

1. 1.

1. Before you stop Linux on your computer, you should configure it so it no longer attempts to start the X server when Linux boots. This is controlled by the **initdefault** line in the /etc/inittab file. Edit this file and change the second field from **5** (Multi-User With X Support) to **3** (Multi-User With No X Support). You could use vi, joe, emacs, or any other suitable text editor to do this job.
2. Perform an orderly shutdown on your system at a safe time. Use the **shutdown -h now** command.
3. Now that the system is off, replace your video card.
4. Start your computer and boot into RHEL. During the boot process, the Red Hat **hal** hardware detection system or the **kudzu** command automatically probes for new hardware. If this probe finds your new video card, you can configure it when prompted.
5. If the automated **hal** or **kudzu** tools fail to find your new hardware, you should use the root account to run the Red Hat Display Settings tool.
6. The Red Hat Display Settings tool should correctly identify your new hardware. You should select the correct amount of display memory (32MB) and the graphics resolutions and color depths you desire. Otherwise, you can configure it manually using the available settings.
7. Test the result. Run the **startx** or **init 5** command to start the Linux GUI.

Lab 2

2. You want to see what happens when there are problems starting the Linux GUI. With RHEL 5, the X.org server is configured by default. The configuration of the X server is stored in the /etc/X11/xorg .conf configuration file. Before Linux starts the X server, it reads this file. To do this lab, you'll want to back up your current /etc/X11/xorg.conf file, delete a line in the file, and then reboot your computer into runlevel 5. You can restore it after the lab is complete.

Answers

2. 1. Back up /etc/X11/xorg.conf to a safe location such as your home directory.

   2. As the root user, delete any active line in the /etc/X11/xorg.conf file.

   3. Open /etc/inittab in your favorite text editor. Look at the line with **initdefault**. Change the number right before this variable from a **3** to a **5** if required.

   4. When you reboot your computer, observe what happens when Linux tries to find the default login display manager. Review the Xorg.0.log file.

   5. Restore your original settings.

   If you are interested in more experiments, try deleting other lines in xorg.conf. Alternatively, try changing ownership of the .Xauthority file for a specific user to root. Log in as that user, run startx, and observe what happens.

Lab 3

3. For this lab, you'll need two Linux computers connected over a network and a shared NFS directory from the local computer. You can use the same directory that you may have used in Chapter 2 to share the RHEL installation files. Start a Secure Shell connection between the two computers. Start the GUI on the local computer, and use the Secure Shell to log in remotely to the other computer.

Once you log in, run the Red Hat root password program from the remote computer. Make changes to the password. When you log out and try to log back into the remote computer, you should be able to confirm that the root password on the remote computer has changed.

Answers

3. For this lab, you'll need two Linux computers connected over a network and a shared NFS directory from the local computer. You can use the same directory that you may have used in Chapter 2 to share the RHEL installation files. You'll start a Secure Shell connection between the two computers. You'll start the GUI on the local computer, and use the Secure Shell to log in remotely to the other computer, with the **-X** or **-Y** switch. You can then see what happens when you start X clients from the remote computer.

   Once you do, run the Red Hat GUI firewall program from the remote computer. Make changes to the firewall, and see what happens. Finally,

   1. On the local computer, start the GUI. If you're currently at the text interface, you can do so with the **startx** command.
   2. Open a command line interface. Assuming you're using the default GNOME desktop, right-click the desktop and click New Terminal from the pop-up menu.
   3. In the new terminal, confirm any currently exported directories with the **showmount -e** command. Based on /etc/exports, select a directory that is set as writable. Use the techniques described in Chapter 10 if required to make it so. You'll be connecting back to one of these directories from your remote computer.
   4. Authorize access from the remote computer. Open the Security Level Configuration tool and allow access through SSH.
   5. Connect to the remote computer using the Secure Shell. Assuming the remote computer is named desktop2, run the following command:

Lab 4

4. In this lab, you'll set up a GUI workstation. It'll start with the **kdm** login manager and automatically start GNOME, open the Firefox Web browser, and start a gnome-terminal session when you boot this Linux computer.

Answers

4. 1. Since you're setting up this workstation for a user, you'll want it to start automatically in the GUI. To do so, open the /etc/inittab file in a text editor, and make sure the **initdefault** variable is set to runlevel 5 as follows:

Chapter 15: Securing Services

Overview

As a Red Hat Enterprise Linux systems manager, you probably wear several hats, one of which is that of security manager. This is especially true if you work for a small company. Even if you work for a large organization with a dedicated network or systems security staff, most of the administrators are probably responsible for other operating systems. You're probably responsible for security policies on your Linux systems.

You may spend very little time thinking about Linux security, or it may turn out to be a full-time job. The level of security you choose to configure depends on many factors, including the purpose of the system and the overall security policies of your company or organization, as well as the size and number of computers in the company.

For example, a Red Hat Enterprise Linux workstation at home does not require as much security as a secure Red Hat Enterprise Linux server that is being used to process credit card orders for a Web site.

Red Hat Enterprise Linux comes with a large and varied assortment of tools for handling security. This includes tools for managing the security on individual Linux computers and tools for managing security for an entire network of systems, both Linux and otherwise. In this chapter, you'll examine some of the tools provided by RHEL for managing security. You'll start out by looking at tools for controlling access to individual Linux host systems, then you'll explore tools for securing networks, and finally, you'll examine the basics of Security Enhanced Linux (SELinux).

Inside the Exam

This chapter is focused on RHCE requirements. As described in the Red Hat Exam Prep guide, RHCEs must be able to

- Configure host-based and user-based security for the service
- Configure SELinux to support the service

for the network services described in the Installation and Configuration portion of the RHCE exam.

These services include HTTP/HTTPS, Samba, NFS, FTP, Web proxy, SMTP, IMAP, IMAPS, POP3, SSH, DNS, and NTP. We've described some security settings in earlier chapters. This chapter looks at several generic security tools that you can use for these services. (For a discussion of Pluggable Authentication Modules, see Chapter 6.)

On the Job

You'll need to know how to protect your computer and network. Sometimes this means you'll turn off, deactivate, or even uninstall a service. Other times, you'll set specific levels of security for different users. You can even regulate the type of traffic coming in, going out, and being transferred through your computer.

Certification Objective 15.01-Using tcp_wrappers to Secure Services

A network is only as secure as the most open system in that network. Although no system can be 100-percent secure, there are certain basic host measures to enhance the security on any given system and, consequently, your network. When devising security measures, you should plan for two types of security violations: user accidents and break-ins.

Accidents happen because users lack adequate training or are unwilling to follow procedures. If security is too burdensome, productivity may suffer, and your users will try to get around your rules. Password rules are sometimes so rigorous, users end up writing their passwords on their desks.

When a cracker breaks in to your system, he or she may be looking for secrets such as credit card information. Others may just want to bring down your system. You can do several things to keep your network secure. Monitor Red Hat errata for the latest issues. Using yum, you can keep your Red Hat system updated with the latest packages.

On the Job

Red Hat is moving away from up2date to yum. RHEL 5 still includes the Red Hat Network registration and management tools such as *rhn_register* and *rhn_check*. If you use Fedora Core 6 or a rebuild such as CentOS to prepare for the RHCE exam, the Red Hat Network is not available to you. Fortunately, knowledge of the Red Hat Network is not part of the publicly listed Red Hat exam requirements.

As you'll see later in this chapter, you can manage your computer's response to certain requests through the /etc/hosts.allow and /etc/hosts.deny files. You can set up protection within the kernel through firewalls based on iptables or **ipchains**. One simple way to promote security is to uninstall as many network access programs as possible.

Security by User or Host

The best way to prevent a cracker from using a service is to remove it completely from your Linux system. However, you may want to keep a service loaded because you're planning to use it in the near future.

You can achieve some measure of security by disabling or removing unused services in the /etc/xinetd.d and /etc/init.d directories. With the services you need, you can block access to specific users, computers, or even networks through the hosts.allow or hosts.deny files in the /etc directory. This system is known as tcp_wrappers, which is enabled by default, and is focused on protecting xinetd services described in Chapter 13.

When a system receives a network request for a service, it passes the request on to tcp_wrappers. This system logs the request and then checks its access rules. If there are no limits on the particular host or IP address, tcp_wrappers passes control back to the service.

The key files are hosts.allow and hosts.deny. The philosophy is fairly straightforward: users and clients listed in hosts.allow are allowed access; users and clients listed in hosts.deny are denied access. As users and/or clients may be listed in both files, the tcp_wrappers system takes the following steps:

1.

It searches /etc/hosts.allow. If tcp_wrappers finds a match, it grants access. No additional searches are required.

2. It searches /etc/hosts.deny. If tcp_wrappers finds a match, it denies access.

3. If the host isn't found in either file, access is automatically granted to the client.

You use the same access control language in both /etc/hosts.allow and /etc/hosts .deny to tell tcp_wrappers which clients to allow or deny. The basic format for commands in each file is as follows:

```
daemon_list : client_list
```

Certification Objective 15.02-Firewalls and Packet Filtering Using netfilter

A firewall sits between your company's internal LAN and an outside network. A firewall can be configured to examine every network packet that passes into or out of your LAN. When configured with appropriate rules, it can filter out those packets that may pose a security risk to your system.

To understand how *packet filtering* works, you have to understand a little bit about how information is sent across networks.

Before you send a message over a network, the message is broken down into smaller units called *packets*. Administrative information, including the type of data, the source address, and destination address, is added to each packet. The packets are reassembled when they reach the destination computer. A firewall examines these administrative fields in each packet to determine whether to allow the packet to pass.

Red Hat Enterprise Linux comes with everything you need to configure a system to be a firewall, including the iptables command.

On the Job

RHEL 5 also includes a firewall command for IPv6 networks, *ip6tables*.

Configuring iptables

The philosophy behind iptables is based on "chains." These are sets of rules applied to each network packet. Each rule does two things: it specifies the conditions a packet must meet to match the rule, and it specifies the action if the packet matches.

The iptables command uses the following basic format:

```
i
```

Certification Objective 15.03-Network Address Translation

Network Address Translation (NAT) lets you hide the IP address of the computers on your network that make a connection to the Internet. NAT replaces the source address with the IP address of the firewall computer, which also serves as a gateway between your network and the Internet. The source address is cached on the gateway, so it knows which computer made the request.

When the firewall receives data such as a Web page, the process is reversed. As the packets pass through the firewall, the originating computer is identified in the cache. The header of each packet is modified accordingly before the packets are sent on their way.

This approach is useful for several reasons. Disguising your internal IP addresses makes it harder for someone to break into your network. NAT allows you to connect computers to the Internet without having to have an official IP address for each computer. This allows you to use the private IP addresses discussed in Chapter 1 on your internal LAN. In the Linux world, this process is known as *IP masquerading*.

IP Masquerading

Red Hat Enterprise Linux supports a variation of NAT called *IP masquerading*. IP masquerading allows you to provide Internet access to multiple computers with a single officially assigned IP address. IP masquerading lets you map multiple internal IP addresses to a single valid external IP address.

Connecting multiple systems to the Internet using IP masquerading is a fairly straightforward process. Your firewall computer will need one network card to connect to your LAN and a second network card for the Internet. This second network card can be a telephone modem, or it can be connected to a cable modem or DSL adapter. This configuration requires the following steps:

1. Assign your official IP address to the network card that is directly connected to the Internet.
2. Assign computers on your LAN one of the private IP addresses described in Chapter 1.
3. Reserve one private IP address for the network card on your firewall that is connected to the LAN.
4. Use iptables to set up IP masquerading.
5. Enable IP forwarding on the firewall computer.
6. Configure the computers on your LAN with the IP address of your firewall computer as their Internet gateway.

Take a careful look at when a message comes from a computer on a LAN, through a firewall, to the Internet. When a computer on your LAN wants a Web page on the Internet, it sends packets to the firewall. The firewall replaces the source IP address on each packet with the firewall's official IP address. It then assigns a new port number to the

packet. The firewall caches the original source IP address and port number.

When a packet comes in from the Internet to the firewall, it should include a port number. If your firewall can match it with the port number assigned to a specific outgoing packet, the process is reversed. The firewall replaces the destination IP address and port number with the internal computer's private IP address and then forwards the packet back to the original client on the LAN.

The next step in the process is to use iptables to enable masquerading. The following command assumes that eth1 represents the network card that is directly connected to the Internet, and that your LAN has a network address of 192.168.0.0/24:

```
#
# iptables -t nat -A POSTROUTING -s 192.168.0.0/24 -o eth1 -j MASQUERADE
```

Certification Objective 15.04-Security Enhanced Linux

Security Enhanced Linux (SELinux) provides one more layer of security. Developed by the US National Security Agency, SELinux makes it more difficult for crackers to use or access any file or service if they break in. SELinux assigns different contexts to each file, known as *subjects*, *objects*, and *actions*.

To see the context of a particular file, run the **ls -Z** command. As an example, review what this command does in Figure 15-3, as it displays security contexts in my /root directory.

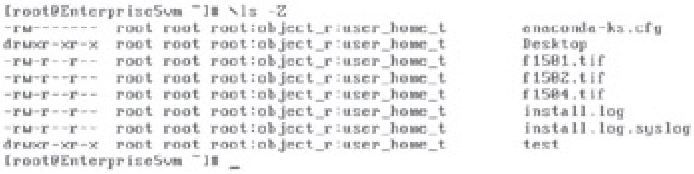
```
[root@Enterprise5vm ~]# \ls -Z
-rw-------  root root root:object_r:user_home_t        anaconda-ks.cfg
drwxr-xr-x  root root root:object_r:user_home_t        Desktop
-rw-r--r--  root root root:object_r:user_home_t        f1501.tif
-rw-r--r--  root root root:object_r:user_home_t        f1502.tif
-rw-r--r--  root root root:object_r:user_home_t        f1504.tif
-rw-r--r--  root root root:object_r:user_home_t        install.log
-rw-r--r--  root root root:object_r:user_home_t        install.log.syslog
drwxr-xr-x  root root root:object_r:user_home_t        test
[root@Enterprise5vm ~]# _
```

Figure 15-3: *ls -Z* output

For this purpose, we'll examine basic configuration tools for SELinux, including the SELinux Management Tool and the Setroubleshoot Browser.

Most SELinux settings are boolean-in other words, they're activated and deactivated by setting them to 1 or 0. Naturally, the booleans are stored in the /selinux/booleans directory. One simple example is **user_ping**, which is normally set to 1, which allows users to run the **ping** command. For a fuller description, see the NSA Guide to Security Policy Configuration using SELinux at www.nsa.gov/selinux/papers/policy/node1.html.

On the Job

If you just want to experiment with SELinux, configure it in Permissive mode. It'll log any violations without stopping anything. It's easy to set up with the Security Level Configuration tool, or you can set *SELINUX=permissive* in /etc/sysconfig/selinux.

SELinux Status

There are three possible statuses for SELinux: **enforcing**, **permissive**, and **disabled**. **enforcing** and **disabled** are self-explanatory. **permissive** means that any SELinux rules that are violated are logged; however, permissive SELinux doesn't stop anything.

If SELinux is active, it protects systems in two ways: in **targeted** or in **strict** mode. The default is **targeted**, and that is what I recommend that you use; it allows you to customize what it protects, and how.

As you'll see shortly, SELinux is easy to configure with the GUI SELinux Management Tool. However, the basics can be easily configured in the /etc/ sysconfig/selinux configuration file. There are three directives in this file, as described in Table 15-3.

Table 15-3: Sample Commands in /etc/sysconfig/selinux

Directive	Description
SELINUX	Basic SELinux status; may be set to **enforcing**, **permissive**, or **disabled**.

SELINUXTYPE	Specifies the level of protection; set to **targeted** by default, where protection is limited to daemons. The alternative is **strict**, which is associated with full SELinux protection.
SETLOCALDEFS	Supports the configuration of local SELinux policies. Set to **0** (disabled) by default.

If you want to change the basic status of SELinux, change the **SELINUX** directive. The next time you reboot, the changes are applied to your system.

Exam Watch

If you have to configure SELinux during your exam, it's no longer possible to do so during the installation process (except to specify *enforcing, permissive,* or *disabled*). If you have to configure SELinux and have to reboot, the process of applying SELinux policies can take several minutes. You won't be able to log in or do anything else during your exam. So plan ahead!

Configuring Manually

SELinux is still relatively new. If you don't understand it well, it may be more efficient to use the SELinux Management Tool to configure SELinux settings. And it's much improved from the GUI SELinux functionality that was part of the Security Level Configuration tool. You can even set SELinux contexts for individual directories from the new tool.

However, as this tool is new, many will believe it's unproven, perhaps until RHEL 6 is released. In Chapter 9, I described how you can set SELinux contexts manually for Apache virtual hosts.

To this end, there are some essential SELinux commands. If you've enabled SELinux, the **ls -Z** command displays the current contexts, as described earlier in Figure 15-3. To see the current status of SELinux, run the **getenforce** command; it returns one of three self-explanatory options: **enforcing**, **permissive**, or **disabled**.

You can change the current SELinux status with the **setenforce** command; the options are straightforward:

```
#
# setenforce enforcing
#
```

Certification Summary

One of the basic functions of a Red Hat Enterprise Linux system administrator is to protect a Linux computer and a network from inside and outside attacks. RHEL includes a variety of tools that can help you establish a secure computing environment.

RHEL includes powerful tools for securing networks from outside attacks. As these tools work in different ways, they provide layers of security. The tcp_wrappers tools use /etc/hosts.allow and /etc/hosts.deny to secure services by host and user. With iptables at your disposal, you can create a firewall that can further protect your RHEL system. The firewall can include IP masquerading to hide the IP addresses of the computers inside your LAN. SELinux uses Access Control Lists to secure individual services in different ways.

The Security Level Configuration tool supports fine-grained configuration of host and network security using iptables. The SELinux Management Tool supports fine-grained control of SELinux targeted policies. The Setroubleshoot Browser can help diagnose SELinux-related problems and even suggest solutions.

Two-Minute Drill

The following are some of the key points from the certification objectives in Chapter 15.

Using tcp_wrappers to Secure Services

? Red Hat Enterprise Linux comes with a package known as tcp_wrappers. This package, which is enabled by default, allows you to limit access to various services.

? You configure the access rules for tcp_wrappers through the /etc/hosts.allow and /etc/hosts.deny configuration files.

? Clients listed in /etc/hosts.allow are allowed access; clients listed in /etc/hosts.deny are denied access.

? Services can also be configured in /etc/hosts.allow and /etc/hosts.deny. Remember to use the actual executable name of the daemon, normally in /usr/sbin, such as **in.tftpd** .

Firewalls and Packet Filtering Using netfilter

? Firewalls can secure an internal network as a packet filter that controls the information that comes in, goes out, and is forwarded through the internal network.

? The current firewall configuration utility is iptables.

? The iptables directives are sets of rules, chained together, that are compared and then applied to each network packet.

? Each rule sets conditions required to match the rule and then specifies the action taken if the packet matches the rule.

? Use the **service iptables save** command to save any chains that you configure in the /etc/sysconfig/iptables configuration file.

Network Address Translation

? NAT modifies the header in packets coming from a LAN. It replaces the source address with the public address of the firewall computer, with a random port number.

? Linux supports a variation of NAT called IP masquerading.

? IP masquerading allows you to provide Internet access to multiple computers with a single officially assigned IP address.

? To enable IP forwarding immediately, type the **echo 1 > /proc/sys/net/ipv4/ip_forward** command. To enable it upon reboot, set **net.ipv4.ip_forward = 0** in /etc/sysctl.conf.

Security Enhanced Linux

? Security Enhanced Linux (SELinux) provides a different level of security. Basic settings are shown in the /etc/sysconfig/selinux file.

? If you're just experimenting with SELinux, configure it in permissive mode.

? SELinux is relatively easy to configure with the SELinux Management Tool.

? Any changes you make with the SELinux Management Tool are reflected in boolean settings in the /selinux/booleans/ directory.

? The Setroubleshoot Browser can help you decipher related errors.

PREV NEXT

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer to many of these questions.

Using tcp_wrappers to Secure Services

1. What happens to a service if you allow the service in /etc/hosts.allow and prohibit it in /etc/ hosts.deny?

---

2. You are using the xinetd program to start services. How could you limit Telnet access to clients on the 192.168.170.0 network? Hint: The telnet daemon, when installed, is in /usr/kerberos/ sbin/telnetd.

Answers

1. If you allow a service in /etc/hosts.allow and prohibit it in /etc/hosts.deny, the service is allowed.

2. You are using the xinetd program to start services. To limit Telnet access to clients on the 192.168.170.0 network, you'd allow access to the network in /etc/hosts.allow and deny it to all others in /etc/hosts.deny. As /usr/kerberos/sbin is in the root user path, you can cite **telnetd** directly and add the following directive to /etc/hosts.allow (remember, CIDR notation doesn't work in these files):

```
telnetd : 192.168.170.0/255.255.255.0
```

Firewalls and Packet Filtering Using netfilter

3. You have recently connected your organization's network to the Internet, and you are a little worried because nothing other than your router is standing between your network and the Internet. You have a spare 400 MHz PC with 256MB of RAM that just happens to have two Ethernet cards. You also have a mixture of systems on your network that includes Macintosh, Windows 2000, and Linux. You also want your system to reject unwanted packets. What might you do to alleviate your concerns?

---

4. Consider the following command:

```
# iptables -A INPUT -s 192.168.77.77 -j REJECT
```

5. What command saves iptables rules?

____

6. Where are iptables rules stored?

____

Answers

3. It's best to create a firewall using the iptables command. The standard Red Hat Security Level Configuration tool creates iptables commands that **REJECT** unwanted packets by default. You can now even support access to the network that can communicate natively with Microsoft and modern Macintosh systems, Samba. All you need is to allow access using the tool.

4. Based on the given command, any connection attempt (including pings) from the 192.168.77.77 system is rejected.

5. The command that saves iptables rules is **iptables-save**.

6. Rules associated with iptables rules are stored in /etc/sysconfig/iptables.

Network Address Translation

7. You are setting up a small office and would like to provide Internet access to a small number of users, but you don't want to pay for a dedicated IP address for each system on the network. What can you do?

____

____

8. What iptables command switch sets up masquerading?

____

Answers

7. To set up a small office while providing Internet access to a small number of users, all you need is one dedicated IP address. The other addresses can be on a private network. Masquerading makes this possible.

8. The iptables command switch that sets up masquerading is **-t nat**.

Security Enhanced Linux

9. What directive activates Security Enhanced Linux in /etc/sysconfig/selinux?

____

10. If you want to let SELinux allow vsFTP service for user home directories, what would you do?

____

11. Where are standard SELinux boolean directives stored?

____

12. If you want to disable SELinux, what would you do?

____

____

Answers

9. The directive in /etc/sysconfig/selinux that activates Security Enhanced Linux is **SELINUX=enabled**.

10. If you want SELinux to allow reading of home directories via an FTP server, activate the Allow Ftp To Read/Write Files In The User Home Directories option. Alternatively, run the **setsebool-P ftp_home_dir 1** command. Additional configuration is required in the vsFTP configuration file, as defined in Chapter 10.

11. Standard SELinux boolean directives are stored in the /selinux/booleans directory.

12. You can disable SELinux in a number of ways. You can do so directly in /etc/sysconfig/selinux by setting **SELINUX=disabled**. You can use the Security Level Configuration tool or even the SELinux Management Tool. You can even add the **selinux=0** directive to the kernel configuration line in your GRUB bootloader. I can even visualize a situation where all these options are used, which would make it more difficult for an RHCE candidate to enable SELinux during an exam.

Lab Questions

Lab 1

1. You want to set up an RHEL computer as a secure Web server. To keep that system secure, you'll want to configure an appropriate firewall and disable any services that you don't need. What should you do?

Answers

1. If you want to set up an RHEL computer as a secure Web server, it's a straightforward process. You'll want to set up a firewall to block all but the most essential ports. This should include TCP/IP ports 80 and 443, which allow outside computers to access your regular and secure Web services.

   The easiest way to set this up is with the Red Hat Security Level Configuration tool, which you can start with the **system-config-securitylevel** command. Once you're in the Red Hat tool, take the following steps:

   1. Enable the firewall. This configures a basic set of firewall rules that prohibits access except for requests that come from inside the firewall.
   2. Scroll down the Trusted Services window. (If you're in the text-based tool, click Customize to open the Firewall Configuration - Customize window.) Activate the WWW (HTTP) option. This allows access from outside the local computer to your regular Web site. Activate the Secure WWW (HTTPS) services as well.
   3. Click OK to exit from the Security Level Configuration tool.
   4. Enter the following command to check your resulting firewall.

Lab 2

2. You want to set up Telnet service on your internal LAN, accessible only to one specific IP address. You want to block access from outside the LAN. Assume that your LAN's network address is 192.168.1.0, and the IP address of the computer that should get access is 192.168.1.33. For the purpose of this lab, feel free to substitute the IP address of a second Linux computer on your network. What do you do?

Answers

2. Several steps are required to set up any xinetd service such as Telnet. You'll need to modify the xinetd Telnet configuration file and set up filtering in one of three ways: in the /etc/xinetd.d/krb5-telnet configuration file, through tcp_wrappers, or with the appropriate firewall commands.

1. First, you want to enable Telnet. Make sure that the krb5-telnet RPM is installed.
2. Activate Telnet. Use the **chkconfig krb5-telnet on** command to revise the /etc/xinetd.d/krb5-telnet configuration script.
3. Edit the /etc/xinetd.d/krb5-telnet configuration file. Add the **only_from = 192.168.1.33** line. (If you have another computer on your network with a private IP address, substitute accordingly in all steps in this lab.)
4. Save the configuration file and reload the xinetd service script with the **service xinetd reload** command. Try accessing Telnet from the local computer. What happens?
5. Try accessing Telnet from the computer with the IP address of 192.168.1.33. What happens? Try again from a different computer on your LAN.
6. Restore the previous /etc/xinetd.d/krb5-telnet configuration file. Don't forget to reload the xinetd service script with the **service xinetd reload** command.
7. Edit /etc/hosts.deny. Add the **telnetd : ALL EXCEPT 192.168.1.33** line.
8. Try accessing Telnet from the computer with the IP address of 192.168.1.33. What happens? Try again from a different computer on your LAN.
9. Restore the previous /etc/hosts.deny file.
10. Save any existing iptables chains. Back up /etc/sysconfig/iptables, if that file currently exists to ~/bak.iptables.
11. Flush current firewall rules with the **iptables -F** command.
12. Block the Telnet port, 23, for all IP addresses except 192.168.1.33 with the **iptables -A INPUT -s ! 192.168.1.33 -p tcp --dport 23 -j DROP** command.
13. Try accessing the Telnet server from the computer with the IP address of 192.168.1.33. What happens? Try again from a different computer on your LAN.
14. Flush current firewall rules with the **iptables -F** command.
15.

Lab 3

3. You want to set up a secure Web server on your corporate LAN that supports inbound requests from your LAN and the Internet, but you do not want any of these requests from the Internet to get into your intranet. What can you do?

   There are three scenarios in this lab. First, assume cost is no object, and there are three computers available-two for firewalls and one for the Web server. Second, assume a cost-conscious situation where you need to configure the firewalls and Web server on the same system. Third, repeat scenarios one and two, with SELinux in **enforcing** mode.

Answers

3. **Scenario 1:** Cost is not an object. This means you can build a DMZ using two firewalls and a separate Web server, all running Linux. You should have the Web server dedicated only to the Web. You configure two more Linux hosts, each with two network cards, and essentially isolate the intranet behind one firewall. You then put the Web server in the middle, placing the second firewall between the Web server and the Internet. You configure the firewall on the intranet with IP masquerading to ensure anonymity for all your intranet hosts.

   Scenario 2: You have one old computer available, and the Web server is a separate computer. Use your one computer as the firewall between you and the Internet and only forward HTTP packets to the Web server IP address directly; use NAT for all intranet requests going out to the Internet for HTTP and FTP. Disallow all other services.

   Scenario 3: Repeat scenarios 1 and 2; configure SELinux in **enforcing** mode and activate the appropriate booleans for the scenarios.

Lab 4

4. You want to work with SELinux, but you are unsure about how it will affect the dozen services that you run from your system. What can you do and what should you monitor to try out SELinux, without affecting any services that are currently running? How can you monitor the process? Test all configured services, and use Setroubleshoot Browser suggestions to configure your system. When you're confident that everything will work, activate SELinux in **enforcing** mode.

Answers

4. The simplest way to experiment with SELinux is to set it to permissive mode. All violations of SELinux are logged in /var/log/messages with the avc label. You can set SELinux to permissive mode with the SELinux Management Tool or by setting **SELINUX=permissive** in /etc/sysconfig/selinux. Open the Setroubleshoot Browser, and try out various services-locally and remotely. Follow any suggestions made by the browser. When you're confident that your configured network services will work with SELinux, set **SELinux=enforcing** in /etc/sysconfig/selinux, reboot, and test configured network services again.

Chapter 16: Troubleshooting

Overview

While you've read about many troubleshooting scenarios throughout this book, it's the troubleshooting part of the Red Hat exams that I believe causes the most "fear and loathing" among Red Hat certification candidates.

Troubleshooting is a mindset based on experience and a systematic way of thinking. Troubleshooting strategies on the Red Hat exams are based on the simplest problems that you can check quickly, moving to more complex problems.

Red Hat has done excellent work addressing some problems that formerly led to unbootable systems. For example, flaws in the /etc/fstab file used to lead to an unbootable system. Now most users would hardly know the difference if this file is missing.

The most important troubleshooting tool is the **linux rescue** environment, which can bypass boot problems, from a missing GRUB boot loader to a missing kernel. In most cases, the first installation CD, booted into the **linux rescue** environment, can detect and mount even damaged installations of RHEL.

This chapter focuses on the Troubleshooting and System Maintenance section of the RHCT and RHCE exams, as defined in the Inside the Exam sidebar. It further focuses on troubleshooting skills, as they evoke more concern than regular system maintenance.

This chapter includes a number of exercises for which you'll need the help of a partner. When you start an exercise, let your partner have your computer and wait until your system begins to reboot. This chapter includes enough exercises to allow you and your partner to take turns working with the system.

Certification Objective 16.01-Troubleshooting Strategies

When you encounter problems, proceed calmly. If you've read this book thoroughly, have the requisite experience, and *do not panic*, you'll usually be able to identify the cause of a problem fairly quickly.

If you can't identify the cause right away, try the simplest solutions first. They take less time and are less likely to sabotage your system.

If you have to go into more detail, remember the seven basic steps of the scientific method (as defined in Wikipedia). They can be applied to the Troubleshooting and System Maintenance portion of your Red Hat exam. If you have experience, you may be able to jump to a solution at any of these steps.

Inside the Exam

Troubleshooting and System Maintenance

As described in the Red Hat Exam Prep guide (www.redhat.com/training/rhce/examprep.html), there are Troubleshooting and System Maintenance requirements for both the RHCT and RHCE exams. To qualify as an RHCE, you need to complete all RHCT requirements during the first hour of the exam. These requirements can fall into the following categories:

- Boot systems into different runlevels for troubleshooting and system maintenance.
- Diagnose and correct misconfigured networking.
- Diagnose and correct hostname resolution problems.
- Configure the X Window System and a desktop environment.
- Add new partitions, filesystems, and swap to existing systems.
- Use standard command line tools to analyze problems and configure system.

To qualify as an RHCE, you also need to complete enough of the RHCE requirements for an overall score of 80%, which can fall in the following categories:

- Use the rescue environment provided by first installation CD.
- Diagnose and correct boot failures arising from boot loader, module, and filesystem errors.
-

- Diagnose and correct problems with network services (see Installation and Configuration below for a list of these services). (The reference is to the Installation and Configuration section of the RHCE exam.)
- Add, remove, and resize logical volumes.
- Diagnose and correct networking services problems where SELinux contexts are interfering with proper operation.

For example, if there are five RHCT problems and five RHCE problems, you'll have to answer all five RHCT problems and three RHCE problems correctly to qualify as an RHCE on this part of the exam.

The network service issues you may encounter may include one or more of the services described throughout this book, including Apache, Samba, NFS, FTP, Squid, sendmail, Postfix, Dovecot, SSH, DNS, and NTP.

---

1. Define the question.

   Understand what happened. Take the error messages you see. If possible, analyze log files for other messages. If you've read this book and run the labs, you may recognize the problem and cause immediately.

2. Gather information and resources.

   Analyze your system. This may require that you check the relevant configuration files to make sure that appropriate services are running and that security or other characteristics of your system are working as they should. If you have experience, you'll often recognize the problem and cause when you see something wrong in these areas.

3. Form a hypothesis.

   If you're still not sure what's wrong, make your best guess. Remember that time is severely limited during the Red Hat exams, so if you can afford it, consider skipping a problem. (To qualify for either the RHCT or RHCE, you're required to solve *all* RHCT-level Troubleshooting and System Maintenance issues.)

4. Perform experiments and collect data.

   Before performing any experiments, back up anything you might change. For example, if you think the problem is with your Samba configuration file, back up your /etc/samba/smb.conf file, in case your hypothesis makes things worse.

5. Analyze data.

   This is essentially identical to step 1. If what you do doesn't solve the problem, you'll need to analyze what went wrong, using error messages and log files as appropriate.

6. Interpret data and draw conclusions that serve as a starting point for new hypotheses.

   In many cases, you'll want to restore what you did from the backup in step 4, repeat steps 2 through 4, and

try again.

7. Publish the results.

   Once you've solved the problem, you'll want to make sure the problem remains solved after rebooting your system. For example, if you've addressed a Samba problem, you'll want to "publish" by making sure the Samba daemon starts the next time your Linux system boots.

Two places where you are likely to make errors that result in an unbootable system are in the boot loader and init configuration files, /boot/grub/grub.conf and /etc/ inittab. For example, identifying the wrong partition as the root partition (/) can lead to a kernel panic. Other configuration errors in /boot/grub/grub.conf can also cause a kernel panic when you boot Linux. Whenever you make changes to these files, the only way to fully test them out is to reboot Linux.

Exam Watch

As a Red Hat Enterprise Linux administrator, you will be expected to know how to fix improperly configured files related to the boot process. For this reason, a substantial portion of the exam is devoted to testing your troubleshooting and analysis skills.

The following scenarios and solutions list some possible problems and solutions that you can have during the boot process, and possible associated solutions. It is far from comprehensive. The solutions that I've listed work on my computer, as I've configured it. There may be (and often is) more than one possible cause. *These solutions may not work for you on your computer or on the Red Hat exams. To know what else to try, use your experience.*

To get the equivalent of more experience, try additional scenarios (remember: never do these things on a production computer). Once you're familiar with the **linux rescue** environment, test these scenarios. These scenarios worked as shown when I tested them on RHEL 5. However, they lead to different errors on RHEL 4 and RHEL 3.

For the first scenario shown, change the name of the grub.conf file so it can't be loaded. Reboot and see what it does on your system. Use the **linux rescue** environment to boot into RHEL and use the noted solution to fix your system.

For the second scenario shown, overwrite the MBR; on a SATA/SCSI drive, you can do so with the following command (substitute **hda** for **sda** if your system uses an IDE/PATA drive):

```
#
# dd if=/dev/zero of=/dev/sda bs=446 count=1
```

Certification Objective 16.02-Required RHCT Troubleshooting Skills

As described in the "Inside the Exam" section at the start of the chapter, the Red Hat Exam Prep guide lists six RHCT-level troubleshooting and system maintenance skills. If you're studying for the RHCE, you must complete all RHCT requirements within the first hour on this part of the exam. You've already read about techniques for booting systems into different runlevels. Now you'll learn about the other RHCT-level items as described in the Red Hat Exam Prep guide.

All the items described in this section have been covered in other chapters. This chapter summarizes important files and commands from those chapters to help focus your thoughts as you move your way through the RHCT-related Troubleshooting and System Maintenance issues.

There are several exercises in this section. Most require the assistance of a partner. Before you perform these exercises, you should have a backup, or a snapshot, of your system, such as that available on a VMware server.

Diagnosing and Correcting Network Problems

To diagnose misconfigured networking, you need to use the commands and analyze the files described in Chapter 7. To check your current network settings, you'll want to run commands such as:

- ifconfig to find the settings of your network card(s)
- **ping** to confirm connectivity to other systems
- **route** to confirm the current routing table

You'll also want to check key files, such as:

- /etc/sysconfig/network to confirm that **NETWORKING=yes**
- /etc/sysconfig/network-scripts/ifcfg-eth0 to confirm defaults for your network card (assuming the default eth0 device for the network card)
- /etc/resolv.conf to confirm connections to DNS servers (which is associated with **PEERDNS=yes** in the aforementioned ifcfg-eth0 configuration file)

Refer to Chapter 7 for more information on these commands and files. There are a lot of details; if you forget something, it may be easier to use a Red Hat utility such as the GUI-based Network Configuration tool. Remember, as long as you don't cheat, it doesn't matter how you solve a problem during the Red Hat exams.

Exercise 16-1: Diagnosing and Correcting Network Problems

For this exercise, you'll need a partner. Have that partner make changes to your system. Let that partner work

privately on your system, until told that the computer is rebooting. Don't look at this lab, until you've solved the problem as created by your partner.

1. Run the ifconfig command and review your current network settings.
2. Back up the configuration file associated with the network card, usually ifcfg-eth0 in the /etc/sysconfig/network-scripts directory. Make sure to back up this file to a non-standard location, in case your partner also backs up any files before changing them.
3. Open up the ifcfg-eth0 file in a text editor.
4. Set **BOOTPROTO=none** if it isn't already done.
5. Set or add an **IPADDR** directive. Make it just a little different from the IP Address setting you saw in the output from ifconfig. Make sure the new address is on a different network; for example, if the original IP Address and network mask was 192.168.0.50 and 255.255.255.0, set **IPADDR=192.168.1.50** and **NETMASK=255.255.255.0**.
6. Reboot your system, and let your partner back at the computer. Tell him or her to try connecting to another system on your network.
7. Tell your partner to back up any files that he or she might change to the home directory.
8. If your partner gives up, restore the original ifcfg-eth0 configuration file to the /etc/sysconfig/network-scripts directory.

Diagnosing and Correcting Hostname Resolution Problems

Hostname resolution is based on the relationship between hostnames such as enterprise5a.example.org and IP addresses such as 192.168.44.66. First, the default hostname is defined in /etc/sysconfig/network, based on the **HOSTNAME** directive. Hostnames are associated with IP addresses in /etc/hosts. If you use a DNS service, you need to make sure that the DNS server's IP address is identified in /etc/resolv.conf. If you use DHCP to get your IP address, it overwrites the DNS server addresses /etc/ resolv.conf, unless **PEERDNS=no** before the **BOOTPROTO=dhcp** directive in /etc/ sysconfig/network-scripts/ifcfg-eth0. Then a command like **dhclient eth0** will acquire the DNS server address(es) from the DHCP server and place them in /etc/resolv.conf.

When there's appropriate routing information, as shown by the **route** command, along with DNS information in /etc/resolv.conf, you can apply the **ping** command to confirm connectivity to the external host of your choice.

Exercise 16-2: Diagnosing and Correcting Hostname Resolution Problems

For this exercise, you'll need a partner. Have your partner make changes to your system. As your partner works to create a network problem for you to solve on your computer, look away until the computer is rebooting.

1. Back up the configuration file associated with the DNS server, /etc/resolv .conf. Back up the /etc/hosts

configuration file. Back up the /etc/host.conf configuration file. Make sure to back up these files to a non-standard location, in case your partner also backs up any files before changing them.

2. Open the /etc/host.conf configuration file in a text editor. If it isn't already as shown, change the directive in this file to:

Certification Objective 16.03-Required RHCE Troubleshooting Skills

At some point in your career as a Red Hat Enterprise Linux administrator, maybe even on the Red Hat exams, you're going to be faced with a system that will not boot. It will be up to you to determine the cause of the problem and implement a fix. Sometimes, the problem may be due to hardware failure: the system in question has a bad power supply or has experienced a hard disk crash.

Quite often, however, the failure of a system to boot can be traced back to the actions of a user: you, the system administrator! When you are editing certain system configuration files, typographical errors can render your system unbootable.

Any time you plan to make any substantial modifications to your system or change key configuration files, back them up first. Then, after making changes, you should actually reboot your system rather than assume that it will boot up the next time you need a reboot. It's much better to encounter problems while you can still remember exactly which changes you made. It is even better if you can go back to a working configuration file.

As described earlier in this chapter, the main tool is the **linux rescue** environment provided by the first installation CD. Per the Exam Prep guide, you also need to know how to *diagnose and correct boot failures arising from bootloader, module, and filesystem errors*; It's broken down into three sections. In addition, some of the key tools to d*iagnose and correct problems with network services*, as described in previous chapters, are summarized here. Key tools are discussed that allow you to *add, remove, and resize logical volumes.* And finally to diagnose and correct networking services problems where SELinux contexts are interfering with proper operation, use the Setroubleshoot browser described in Chapter 15.

Troubleshooting the Boot Loader

The boot loader associated with Red Hat Enterprise Linux 5 is GRUB. For an extensive discussion, see Chapter 3. It can help you to know how to:

- Associate the root directive with the partition with the /boot directory.
- Boot into the desired, non-default runlevel.
- Access the GRUB command line.
- Test different GRUB commands.
- Use command completion to find and use the exact names of your kernel and initial RAM disk.

While it isn't necessary to know all these skills, they can help you diagnose problems more quickly during your exam.
Exercise 16-6: Troubleshooting the Boot Loader

[redacted]is exercise, you'll need a partner. Have your partner make changes to your system. As your partner works to

create a network problem for you to solve on your computer, look away until the computer is rebooting.

It's most helpful if you have a VMware snapshot of your RHEL system. Problems like those created in this exercise have caused administrators to mess up their systems in other ways. You'll also need the first RHEL installation CD.

1. Back up the configuration file associated with the boot loader, /boot/grub/grub.conf. Make sure to back up this file to a non-standard location, in case your partner also backs up any files before changing them.

2. Open the /boot/grub/grub.conf configuration file in a text editor. Focus on the kernel command line, which might look like one of the following:

Certification Summary

One of the most valuable skills as a systems administrator is knowing how to troubleshoot a system. Understanding the scientific method is the first skill. Knowing how to boot into different runlevels and how to use the **linux rescue** mode from the first installation CD are other key skills.

The Troubleshooting and System Maintenance portion of the Red Hat exams includes RHCT and RHCE skills. If you're taking the RHCE exam, you'll need to pass both sections to qualify for the RHCE. The RHCT section requires that you know how to diagnose and correct misconfigured networking, diagnose and correct hostname resolution problems, configure the X Window System and a desktop environment, add new partitions and filesystems, and swap to existing systems. You also need to know how to use standard command line tools and configure your system. All of these skills were covered in this book and are summarized in this chapter.

RHCE candidates also have their own set of challenges during the Troubleshooting and System Maintenance portion of that exam. You'll need to know how to diagnose and correct boot failures arising from boot loader, module, and filesystem errors, as well as add, remove, and resize logical volumes. You'll also need to know how to diagnose and correct problems with any of the network services described throughout the book as well as related issues where SELinux is interfering with proper operation of these services.

Two-Minute Drill

Here are some of the key points from the certification objectives in Chapter 16. **Troubleshooting Strategies**

? During the Troubleshooting and System Maintenance exam, use your experience. You may have seen the problem before.

? If you aren't sure about a problem, try the simplest solution first. q If all else fails, use the scientific method.

? Know how to boot Linux into different runlevels; it can help you bypass many problems and boot a system.

? RHCE candidates should also know how to boot Linux from the first installation CD, using **linux rescue** mode.

Required RHCT Troubleshooting Skills

? Misconfigured networking can be diagnosed with key commands such as ifconfig and **ping**. It can be corrected in related configuration files or with the Red Hat Network Configuration tool.

? Misconfigured hostname resolution problems relate to connections to /etc/ hosts or DNS servers.

? Configure the X Window System using the Red Hat Display Configuration tool, commands such as **Xorg -configure**, or by directly editing /etc/X11/xorg.conf.

? Configure a desktop environment and login manager with key configuration files such as /etc/X11/prefdm and /etc/X11/xinit/xinitrc.

Required RHCE Troubleshooting Skills

? Knowing how to diagnose and correct boot failures means knowing the GRUB configuration file and command line.

? Understand how to manage boot modules with the correct initial RAM disk.

? If there are filesystem errors, they may show up during the boot process and require that you use commands such as fsck to diagnose and solve them.

? Adding, removing, and resizing LVs require skills described in detail in Chapter 8.

? Know how to diagnose and solve service issues as described throughout this book.

? Use the Setroubleshoot browser to diagnose network service problems related to SELinux.

Self Test

The following questions will help you measure your understanding of the material presented in this chapter. As no multiple choice questions appear on the Red Hat exams, no multiple choice questions appear in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams. There may be more than one answer for many of these questions.

As this chapter contains references to many other chapters in this book, some of the questions here require knowledge from reading other chapters.

Troubleshooting Strategies

1. What runlevel options are there if you do not want to boot into runlevels 0, 2, 3, 4, 5, or 6? ?

2. If **linux rescue** mode successfully finds an existing RHEL installation, where is it mounted? ?

Answers

1. If you do not want to boot into runlevels 0, 2, 3, 4, 5, or 6, you can boot into runlevels 1, s, or emergency; you can also boot with the **init=/bin/sh** command appended to the end of the kernel command line.

2. If **linux rescue** mode successfully finds an existing RHEL installation, it is mounted on the /mnt/sysimage directory.

Required RHCT Troubleshooting Skills

3. If you want to set network card eth0 to acquire IP address information from a DHCP server, what directive would you use in /etc/sysconfig/network-scripts/ifcfg-eth0? ?

4. What files are associated with hostname configuration during the boot process? ?

5. If you want to start from scratch and create a new xorg.conf configuration file, what command would you use?

____

6. If you want to set up a default application in the X Window that will work for both GNOME and KDE, what file or directory would you change?

____

7. What command formats /dev/sda5 to the ext3 filesystem?

____

Answers

[3.](#) If you want to set network card eth0 to acquire IP address information from a DHCP server, you would use the **BOOTPROTO=dhcp** directive in /etc/sysconfig/network-scripts/ifcfg-eth0.

[4.](#) The files that determine the hostname when you boot Linux are /etc/hosts and /etc/sysconfig/ network.

[5.](#) If you want to start from scratch and create a new xorg.conf configuration file, you could use the **system-config-display** command to create it in the appropriate /etc/X11 directory. Alternatively, if you use the **Xorg -configure** command, you'll have to copy xorg.conf from the local directory to /etc/X11.

[6.](#) If you want to set up a default application in the X Window that will work for both GNOME and KDE, you would change /etc/X11/xinit/xinitrc or files in the /etc/X11/xinit/xinitrc.d/ directory.

[7.](#) The command that formats /dev/sda5 to the ext3 filesystem is **mkfs.ext3 /dev/sda5**.

Required RHCE Troubleshooting Skills

8. Other than possibly the version number of the kernel, what's wrong with the following line from /boot/grub/grub.conf?

```
kernel /vmlinuz-2.6.18-8.el5 ro root=LABEL=/dev/hda2 rhgb quiet
```

9. What's wrong with the following line from /boot/grub/grub.conf?

```
initrd /initrd-2.6.18-8.el5
```

10. If you have a problem with the Samba configuration file and suspect a syntax error, what command would you use?

____

11. If you have a problem with the Apache configuration file and suspect a syntax error with virtual hosts, what command would you use?

____

12. What command increases the size of an existing LV, /dev/thisvg/lvol1, by 1000MB, using space available from a properly configured /dev/sda10 device?

?

____

Answers

8. The line from /boot/grub/grub.conf is missing the forward slash in front of /vmlinuz-2.6.18-8.el5; it's the /boot directory as defined by the **root(hd0,0)** (or similar) directive. It should read as follows:

```
k
kernel /vmlinuz-2.6.18-8.el5 ro root=LABEL=/dev/hda2 rhgb quiet
```

9. The line from /boot/grub/grub.conf is missing the .img extension at the end of the initial RAM disk file; it should read as follows:

```
i
initrd /initrd-2.6.18-8.el5.img
```

10. If you have a problem with the Samba configuration file and suspect a syntax error, you would use the **testparm** command.

11. If you have a problem with the Apache configuration file and suspect a syntax error with virtual hosts, you could use either of the following commands:

```
#
# httpd -S
#
```

12. The command that increases the size of an existing LV, /dev/thisvg/lvol1, by 1000MB, using space available from a properly configured /dev/sda10 device, is shown here:

```
#
# lvextend -L+1000 /dev/thisvg/lvol1 /dev/sda10
```

Lab Questions

Lab 1

1. For this lab, you'll need a partner. The first steps will set up the lab for that person. To prepare this lab, take the following steps:
    1. Log in as the root user.
    2. Open /etc/inittab in a text editor.
    3. Change the default runlevel to 0.
    4. Save your changes.
    5. Power down the computer. Pass the system to your partner. Tell your partner that the problem should appear when the computer boots into Linux.

       ****
    6. Now that your partner has set up this system for you, power on the computer and boot into Linux. What happens? What do you see? What can you do?

Answers

1. To solve this problem, you need to observe what happens when you boot this system. While not required, the first step you should take is to use the GRUB boot menu to boot into a specific runlevel. If you have any experience with Linux, you should know that an immediate shutdown after Linux goes through the boot process is associated with runlevel 0.

Lab 2

2. For this lab, let your partner set up your computer for you.

****

Now you'll set up this system for your partner using the following steps.

1. If you've configured your RHEL system on VMware, make sure you have a current snapshot. It's invaluable if your partner is unable to solve this problem.
2. Boot into your partner's system normally.
3. Open the /etc/fstab configuration file. There should be a line associated with the /boot directory, similar to:

Answers

2. In this lab, you should know almost immediately that there's a problem with the **LABEL** associated with the /boot directory, with an error message similar to:

```
fsck.ext3: Unable to resolve 'LABEL=/boot'
```

Lab 3

3. For this exercise, use a test computer. Do not use a production computer. Do not use a computer with any data might be important to you. If something goes wrong, and you are unable to restore from a backup, you may need to reinstall Linux. This exercise assumes that you're using the default Red Hat Enterprise Linux boot loader, GRUB.

Navigate to the /boot directory. Change the name of the initrd-*versionnumber*.img file. Make sure it's something easy to remember such as initrd-*versionnumber*.bak. Reboot Linux. As GRUB goes through the boot sequence, it will probably stop when it can't find your initial RAM Disk (initrd) file, with a file not found message.

Now that your boot loader isn't working, what do you do? Can you try to start Linux in single-user mode?

Answers

3. As you practice learning about Linux for the RHCE exam, it's important to know how GRUB works. By default, it requires an initial RAM disk file, initrd-*versionnumber*.img. If GRUB can't find this file, it'll give you a file not found error. Since your computer does not boot, you'll need to boot with a rescue disc before you can fix the initrd file. Remember to make sure that the filename matches the name shown in /boot/grub/grub.conf *exactly*.

You can repeat this process with the vmlinuz file or the root directive in grub.conf. Make sure to have backups of key files so you can restore your original configuration. When you repeat this process, what happens after you select a kernel from the GRUB menu? Do you see a different error? Is it associated with a different file?

Understanding these answers can help you learn to use GRUB messages to more precisely diagnose specific problems with Linux.

Lab 4

4. In this lab, you'll create new PEs and use them to increase the size of a configured LV. You're doing this for the LV used by the /var directory. Because of the increasing demands of your Web site, you need more room for the /var directory for your Web site data. Assume your /etc/fstab configuration file includes the following line:

```
/
/dev/Volume00/LogVol00 /var ext3 defaults 1 2
```

Answers

4. If you've just added the new hard drive, you'll need to set up partitions or use the entire hard drive for PEs. Based on the premises of the lab, you have the entire SCSI /dev/sdg hard drive available, so you can just allocate this entire hard drive as PEs with the following command:

```
#
# pvcreate /dev/sdg
```

PREV NEXT

Appendix A: Sample Exam 1

Overview

The following questions will help you measure your understanding of the material presented in this book.

As discussed in the introduction, the RHCE exam consists of two different, equally weighted sections: Troubleshooting and System Maintenance (2.5 hours) and Installation and Configuration (3.0 hours). There are RHCE and RHCT components to each section. To earn the RHCE, you need to meet all of the following requirements. To earn the RHCT, you need to meet the noted RHCT requirements.

On the Troubleshooting and System Maintenance section:

- Successfully complete all five RHCT-level problems within the first hour.
- Answer enough RHCE requirements (three of five) to earn an overall section score of 80 (of 100).

On the Installation and Configuration section:

- Earn a score of 70 or higher on the RHCT components.
- Address enough issues correctly to get a grade of 70 on the RHCE components.

Both exams are "closed book." However, you are allowed to use any documentation, such as man pages, that you can find on the Red Hat Enterprise Linux computer. You are allowed a pen or pencil and paper to make any notes that you might need.

If you pass the RHCT components of both sections, you'll qualify as a Red Hat Certified Technician. If you pass the RHCT and RHCE components of both exams, you will also become a Red Hat Certified Engineer.

In most cases, there is no one solution, no single method to solve a problem or install a service. There are a nearly infinite number of options with Linux, so I can't cover all possible scenarios.

Troubleshooting and System Maintenance

For this sample section, you need a computer that you're willing to dedicate for experimental purposes. Actual troubleshooting questions require the installation of the latest version of Red Hat Enterprise Linux, configured with a specific problem. The exam conditions would delete any and all data that you have on that computer. One option is a virtual machine solution such as that provided by VMware; as of this writing, no-cost subscriptions for VMware Server are available at www.vmware.com. I used VMware, using the tips that I describe in "Studying with a Virtual Machine" in the Online Learning Center (http://highered.mhhe.com/sites/0072264543), to write much of this book. Another alternative is Xen, also described in "Studying with a Virtual Machine."

If possible, get a friend, fellow student, or colleague to help set up the exercises for this exam. That's the best way to simulate real-world conditions. As shown in the RHCE and RHCT Exam Prep guide at www.redhat.com/training/rhce/examprep.html, you may have to boot into different runlevels, solve networking and host name problems, configure the GUI, add partitions of various types, and use standard command line tools to analyze and configure your system. You can use any documentation that you can find on your Red Hat Enterprise Linux computer; however, you're not allowed to reinstall Linux to address these problems.

You can't pass either exam unless you solve all of the RHCT-level troubleshooting problems. On the other hand, you need to budget your time judiciously; if you can't solve one RHCE-level problem, you may want to give up and move on to the next problem. But you may not be able to go back. You may be able to debug the next problem in just a few minutes. Even if you have time left over at the end of the section, you may not be able to go back and will not get any credit for any problems that you abandon.

These are not actual questions, but exercises consistent with the guidelines in the Red Hat Exam Prep guide. As exercises, they have no answers per se; however, they include a lot of information that can help you as a Linux administrator. However, I've set them up in a format that can allow someone else to set up exercises similar to what I'm guessing you *might* see on the Red Hat exams.

Even for these exercises, *do not use a production computer.* Some or all of these exercises are designed to make Linux unbootable. If you're unable to recover from the steps documented in these exercises, you may need to reinstall Red Hat Enterprise Linux. Saving any data that you have on your computer at that point may not be possible.

Troubleshooting and System Maintenance Exercise: RHCT Components

There are five problems on the RHCT part of this section. You have to reconfigure your computer to address all five problems appropriately, within the first hour of your exam. In this exercise, you'll set up five different problems on the same computer or virtual machine.

Ideally, you'll have a friend or classmate who can help you prepare your computer. Assume that you have a network server with the RHEL installation files. When your friend or classmate installs RHEL on this text computer, assume that you're on a network with an IP address of 10.20.30.0 (you can substitute a different network IP address if you want). Include the GNOME Desktop Environment in the installation. Then have your friend or classmate take the following steps:

1. Configure an RHEL system with some empty space on a hard drive. For the purpose of this section, I'll assume that you've configured a system with at least 5000MB of free space.
2.

If you have a DHCP server on your network, disconnect or deactivate it.

3.

Configure an IP address for your network card of 10.20.30.40, with a network mask of 255.255.255.0. You can use a tool such as **system-config-network**. Once saved, assign an IP address in /etc/hosts of 10.20.30.50 to the local host name; for example, on my /etc/hosts, I've configured:

Installation and Configuration

The Installation and Configuration section is the second part of the Red Hat exam. If you're taking the RHCT, you're allowed 2.0 hours; if you're taking the RHCE, you're allowed 3.0 hours to install and configure Red Hat Enterprise Linux. You may get partial credit on some of these problems. You have access to the Red Hat Enterprise Linux installation files through a network server (the Exam Prep guide specifies a network installation). Once Red Hat Enterprise Linux is installed, you also have access to the man pages as well as any other documentation that you may have installed.

If you're studying for the RHCT, you can limit your focus to the RHCT-level skills; if you're studying for the RHCE, you'll need to complete all RHCT- and RHCE-level skills, in the time allotted.

No specific techniques or commands are required. Any reasonable technique is allowed if it gets you to the objective. For example, if you need to limit access to a specific service, you can use iptables, /etc/hosts.deny, or even SELinux. As long as it does the job, the configuration can get you full credit for that part of the exam.

You may need to limit access to network servers to specific users or other computers. However, this is a certification exam. Do not expect to have *physical* access to any other computer to test your settings. You will not have access to any outside networks such as the Internet.

If you're going for your RHCT, you'll need a grade of at least 70 percent in the RHCE-level skills. If you're going for your RHCE, you'll need a score of 70 percent on both sections.

Exam Watch

Read the entire Installation and Configuration exam before you finish installing RHEL. It's easier to configure RAID and logical volumes during the installation process. It can save time to install required servers during the installation process. And remember, you can start configuring RHEL through the CTRL-ALT-F2 console even while packages are being installed.

If you're preparing for the RHCT exam, you can ignore the RHCE issues. If you're preparing for the RHCE exam, you'll want to address *all* requirements in this section. Remember, the RHCE is inclusive of the RHCT.

Server Installation Problem: RHCT-Level Skills

Install Red Hat Enterprise Linux. The following conditions specify a network server, configured with some very specific partitions. You'll also need to limit access to some or all of your network servers to specific users, computers, entire networks, or more.

Install Linux over a network connection with the partitions shown in Table A-1. The sizes shown are minimums. Use a reasonable size for the swap partition.

Table A-1: Available Red Hat Enterprise Linux Kernels (and Related Packages)

Filesystem	Size
/boot	100MB
/	4000MB

/home	1000MB
/var	1000MB

You'll want a RAID 6 array for the /home directory where your users can store at least 1000MB of data. Assume this computer has an IP address of 10.11.12.13 on the 10.11.12.0/255.255.255.0 network.

Configure the following users for RHEL: nancy, randy, donna, and mike. Make nancy and randy part of a group named angels. Create a /home/angels directory and allow them to share files without having to change permissions or ownership on any file they put in this directory. Do not give donna or mike read privileges on this directory. Configure quotas for donna and mike to limit the space available in their home directories to 100MB.

Set up Access Control Lists on the /home directory partition. Set up a project .test file in user mike's home directory. Configure ACLs on project.test to allow user donna to read this file.

Make **kdm** the default window manager. Make sure users are directed to the graphical login interface when RHEL starts on this computer. Configure a connection to an LDAP client, on the vtc.com domain, on IP address 10.11.12.15.

Set up a job to delete all of the regular files in the /home/mike directory on the second day of every month at 3:50 A.M. Configure the automounter to connect to the NFS installation source on the /var/ftp/pub directory from IP address 192.168.0.50 (substitute the directory and IP address from your own network accordingly). Connect to and configure a remote CUPS printer; make it the default for this computer.

Install The GIMP after installation. Install the later version of the kernel that's available from the network installation source. Set up another GRUB stanza to boot your system in runlevel 1.

Finally, allow the local system to accept source routing. While normally disabled, it's often associated with systems configured as routers.

Server Installation Problem: RHCE-Level Skills

In this part of the exam, you'll configure a number of different servers on the RHEL computer.

When you install, configure a logical volume, dedicated to the /var directory. Enable ssh logins, and limit access to the local network. Configure Samba to share the /home/angels directory with the users specified earlier. Configure a vsFTP server. Limit access to computers on the LAN. Support access from users over FTP from one other system on the LAN to their home directories. Set up a local NTP server, accessible to other workstations on the LAN.

Configure an NFS server to share the files in /tmp only with users on the LAN. Configure Apache to serve a homepage.html page from within the /var/www/html directory. Do not limit access to the computers on the LAN. Add a secure.html page for connections to a secure Web server. Set up a proxy server that can be used by other computers on your LAN. Make sure SELinux settings allow these options.

Set up a Kickstart file, ready to use with a boot CD. Copy it to a USB key in that system's top-level directory. Configure a local caching nameserver. Add a non-secure POP3 server for the local network; do not support IMAP or secure protocols on that server.

When you reboot your computer, all of the services that you've created and settings that you've made should be enabled automatically.

Installation Discussion

Since there is no one way to set up a Red Hat Enterprise Linux configuration, there is no one right answer for the listed requirements. But there are some general things to remember. You need to make sure your changes work after a reboot. If you're going for the RHCE, you'll need to make sure that the services that you set up are active at the appropriate runlevels. For example, if you're configuring Apache, it should be active for at least runlevels 3 and 5.

First, examine the RHCT-level skills. You can set up the required partitions through Disk Druid during the RHEL installation process, or with the fdisk or **parted** utilities after RHEL is installed. Remember, you need (at least) four partitions for a RAID 6 array. If you're creating the array after installation, use the mdadm command.

Remember your CIDR notation; 10.11.12.0/24 specifies a subnet mask of 255.255.255.0. Use the SGID bit and assign 770 permissions on /home/angels, and make sure to assign group ownership of that directory to angels. Before you can configure quotas, you'll need to remount /home with at least the **usrquota** setting, and add it to /etc/fstab. Create appropriate quota configuration files with quotacheck **-cuvm** (or reboot); and then activate quotas with quotaon; configure quotas for users donna and mike, using the edquota command.

Before configuring ACLs, you need to set up the applicable partition with **acl** settings. You should do so in /etc/fstab. To make it work before a reboot, remount the /home directory partition with the **mount -o remount,acl /dev/partitionnumber /home** command. If it works, you'll be able to confirm with the mount command by itself.

Make your default window manager through the /etc/X11/prefdm file; in this case, you can set **preferred=kdm**. Remember, the GUI login is associated with runlevel 5 in /etc/inittab. Setting up a connection to an LDAP server means understanding the distinguished names associated with the vtc.com domain, where **dc=vtc** and **dc=com**, and the LDAP server is on IP address 10.11.12.15.

Setting up a job to delete files in any specific directory on a periodic basis is a job for the cron daemon. You can create your own job with crontab, or you can set it up through a script similar to **/etc/cron.daily/tmpwatch**. For example, I ran **crontab -e** as user michael and added the following:

```
S
SHELL=/bin/bash
```

Appendix B: Sample Exam 2

Overview

The following questions will help you measure your understanding of the material presented in this book.

As discussed in the introduction, the RHCE exam consists of two different, equally weighted exams: Troubleshooting and System Maintenance (2.5 hours) and Installation and Configuration (3.0 hours). There are RHCE and RHCT components on each exam. To earn the RHCE, you need to meet the following requirements.

On the Troubleshooting and System Maintenance exam, you need to

- Successfully complete all five RHCT-level problems within the first hour.
- Answer enough RHCE requirements (three of five) to earn an overall section score of 80 (of 100).

On the Installation and Configuration section, you need to

- Earn a score of 70 or higher on the RHCT components.
- Address enough issues correctly to get a grade of 70 on the RHCE components.

Both exams are "closed book." However, you are allowed to use any documentation that you can find on the Red Hat Enterprise Linux computer, such as man pages. You are allowed a pen or pencil and paper to make any notes that you might need.

If you pass the RHCT components of both exams, you'll qualify as a Red Hat Certified Technician. If you pass the RHCT and RHCE components of both exams, you will also become a Red Hat Certified Engineer.

In most cases, there is no one solution, no single method to solve a problem or install a service. There are a nearly infinite number of options with Linux, so I can't cover all possible scenarios.

Troubleshooting and System Maintenance

For this sample section, you need a computer that you're willing to dedicate for experimental purposes. Actual troubleshooting questions require the installation of the latest version of Red Hat Enterprise Linux, configured with a specific problem. The exam conditions would delete any and all data that you have on that computer. One option is a virtual machine solution such as that provided by VMware; as of this writing, no-cost subscriptions for VMware Server are available at www.vmware.com. I used VMware, using the tips that I describe in "Studying with a Virtual Machine" in the Online Learning Center (http://highered.mhhe.com/sites/0072264543), to write much of this book. Another alternative is Xen, also described in "Studying with a Virtual Machine."

If possible, get a friend, fellow student, or colleague to help set up the exercises for this exam. That's the best way to simulate real-world conditions. As shown in the RHCE Exam Prep guide at www.redhat.com/training/rhce/examprep.html, you may have to boot into different runlevels, solve networking and host name problems, configure the GUI, add partitions of various types, and use standard command line tools to analyze and configure your system. You can use any documentation that you can find on your Red Hat Enterprise Linux computer; however, you're not allowed to reinstall Linux to address these problems.

You can't pass either exam unless you solve all of the RHCT-level troubleshooting problems within the first hour. On the other hand, you need to budget your time judiciously; if you can't solve one RHCE-level problem, you may want to give up and move on to the next problem. But you may not be able to go back. You may be able to debug the next problem in just a few minutes, but even if you have time left over at the end of the section, you may not be able to go back to skipped problems and will not get any credit for the problems that you have abandoned.

These are not actual questions, but exercises consistent with the guidelines in the Red Hat Exam Prep guide. As exercises, they have no answers per se; however, they include a lot of information that can help you as a Linux administrator. However, I've set them up in a format that can allow someone else to set up exercises similar to what I'm guessing you *might* see on the Red Hat exams.

Even for these exercises, *do not use a production computer.* Some or all of these exercises are designed to make Linux unbootable. If you're unable to recover from the steps documented in these exercises, you may need to reinstall Red Hat Enterprise Linux. Saving any data that you have on your computer at that point may not be possible.

Troubleshooting and System Maintenance Exam: RHCT Components

There are five problems on the RHCT part of this section. You have to reconfigure your computer to address all five problems appropriately, within the first hour of your exam. In this exercise, you'll set up five different problems on the same computer or virtual machine.

Ideally, you'll have a friend or classmate who can help you prepare your computer. Assume that you have a network server with the RHEL installation files. When you install RHEL on this text computer, assume that you're on a network with an IP address of 172.16.32.0 (you can substitute a different network IP address if you desire). Include the K Desktop Environment in the installation. Then have your colleague take the following steps:

1. Install RHEL with at least 256MB of space that is unallocated to any partition. Create a /data directory on a logical volume.
2. Make sure RHEL boots from the command line, by setting the **id** directive in /etc/inittab to runlevel 1.
3. Set up Firefox to open automatically in the default GNOME Desktop Environment, using the Sessions manager (which you can open with the **gnome-session-properties** command in the GUI). Remember, this is user dependent, so you'll have to specify the user whose GUI the candidate is to configure.
4. Open /etc/X11/xorg.conf; create a typo in the **Screen** directive as defined in the first stanza. In my configuration, I changed

Troubleshooting and System Maintenance: RHCE Components

You have to solve three of the five simulated RHCE-level problems correctly. Your time starts after you've answered *all* of the simulated RHCT problems correctly; thus, you'll have at least 1.5 hours for this part of the exam. Ideally, you'll have a friend or classmate help you prepare your computer. Assume that you have access to a network server with the RHEL installation files.

You may want to have your friend or classmate prepare additional questions. Some boot exercises can be easily created from the Scenario and Solution list in Chapter 16. Other exercises on network services can be created based on what you've learned in Chapters 7 and 9?15. Exercises related to adding, removing, and resizing logical volumes can be developed from the information in Chapter 8.

In these exercises, you'll be working with an RHEL computer with some key files messed up. Ideally, you'll have a friend or classmate help you prepare your computer. Assume that you have a network server with the RHEL installation files. This should be on a system with some free unpartitioned space on the hard drive, and one filesystem, /tmp, on a logical volume.

You'll also need a computer on which you can boot directly from your CD drive, so you can use the first Red Hat installation CD. (If you can boot from the USB, you can substitute a specially prepared boot USB drive described in Chapter 2.) Then walk away, and have the colleague who is preparing the exam take the following steps:

1. Copy and back up the /etc/inittab configuration file. One possible name is /etc/bak.inittab.
2. Open /etc/inittab file in a text editor.
3. Change the *x* in the **id:*x*:initdefault** line to 4.
4. Comment out the **l4:4:wait:/etc/rc.d/rc 4** line in /etc/inittab. Save your changes.
5. Activate SELinux in enforcing mode.
6. Change the kernel directive in the GRUB configuration file (/boot/grub/grub.conf), to point the root directive as shown:

Troubleshooting and System Maintenance Discussion

You shouldn't read this discussion until you've had a chance to try out the problems. I describe one possible solution to each problem. The solutions that you come up with can vary. The method you use doesn't matter; the result is what counts. As of this writing, in the Troubleshooting and System Maintenance section, you can't reinstall RHEL on the target computer and you can't go to the Internet for help during the exam.

Troubleshooting and System Maintenance Exam: RHCT-Level Problems

To pass either exam, you need to solve all five RHCT-level Troubleshooting and System Maintenance problems. In this simulation, you'll need to

- Allocate half of the free space on the current drive to a partition dedicated to the /test directory.
- Make Linux boot into the GUI.
- Address any problems you see the next time you boot Linux.
- Configure the GNOME desktop to open a command line window for the selected user when you boot.
- Assign a host name of enterprise5 to this system, and configure the network card to connect with DHCP.

Address the problems that you've been given. In most cases, there are other ways to solve these problems. When you boot this system, the first thing you'll see is this:

```
s
s ħ3 .1#
```

Installation and Configuration

The Installation and Configuration section is the second part of the Red Hat exam. If you're taking the RHCT, you're allowed 2.0 hours; if you're taking the RHCE, you're allowed 3.0 hours to install and configure Red Hat Enterprise Linux. You may get partial credit on some of these problems. You'll have access to the Red Hat Enterprise Linux installation files through a network server (the Exam Prep guide specifies a network installation). Once Red Hat Enterprise Linux is installed, you'll also have access to the man pages as well as any other documentation that you may have installed.

If you're studying for the RHCT, you can limit your focus to the RHCT-level skills; if you're studying for the RHCE, you'll need to complete all RHCT- and RHCE-level skills, in the time allotted.

No specific techniques or commands are required. Any reasonable technique is allowed if it gets you to the objective. For example, if you need to limit access to a specific service, you can use iptables, /etc/hosts.deny, or even SELinux. As long as it does the job, the configuration can get you full credit for that part of the exam.

You may need to limit access to network servers to specific users or other computers. However, this is a certification exam. Do not expect to have *physical* access to any other computer to test your settings. You will not have access to any outside networks such as the Internet.

If you're going for your RHCT, you'll need a grade of at least 70 percent in the RHCT-level skills. If you're going for your RHCE, you'll need a score of 70 percent on both sections.

Exam Watch

Read the entire Installation and Configuration exam before you finish installing RHEL. It's easier to configure RAID and logical volumes during the installation process. It can save time to install required servers during the installation process. And remember that you can start configuring RHEL through the CTRL-ALT-F2 console even while packages are being installed.

Most of you will find it difficult to complete this exercise within 3 hours. I've deliberately added extra difficulty to this second sample exam, which will hopefully ease your required effort during the actual RHCE exam. I've also split up this sample exam into RHCT- and RHCE-level skills for your convenience; it may not represent what you'll actually see on the exam. (I'm not allowed to tell you about it.)

If you're preparing for the RHCT exam, you can ignore the RHCE issues. If you're preparing for the RHCE exam, you'll have to meet *all* requirements in this section. Remember that the RHCE is inclusive of the RHCT.

Once you've mastered the skills in this book, try other variations. Practice with different scenarios until you become comfortable with the scenarios described in this book, as well as in the Red Hat Exam Prep guide.

Server Installation Problem: RHCT-Level Skills

Install Red Hat Enterprise Linux. The following conditions specify a connection to network servers, configured with some very specific partitions. Assume this computer gets its IP addressing information from a DHCP server. Let users start at a virtual console.

Install Linux with the partitions shown in Table B-1. The sizes shown are minimums. Make sure that the /home directory is configured in a RAID 5 software array with no spare partitions. Leave 1000MB of free, unallocated space on the hard drive. If your system has less available hard drive space, some adjustments may be possible, such

as reducing the amount of space allocated to / and /usr to 2000MB each.

Table B-1: Required Partitions

Filesystem	Size
/boot	100MB
/	4000MB
/home	1000MB
/var	1000MB
/usr	4000MB

Once RHEL is installed, you'll also want to configure the following:

- A connection to a local printer
- Active networking only during working hours (8:00 A.M to 5:00 P.M.)
- An NIS client, on the biglan NIS domain
- The automounter, configured to read the CD on the /misc/cd directory
- Support for IP forwarding, as this computer may be a router in the future
- Installation of the system-config-boot RPM
- A Linux kernel, upgraded to the latest requirements

Configure a cross-functional group of users: avionics, vendor, seats, and galleys. Set them up as a group named pcplane. Create a /home/pcplane directory and allow them to share files without having to change permissions or ownership on any file they put in this directory. Do not give vendor read privileges on this directory. Limit each of these users to 100MB of files in this directory. Make it possible to create ACLs on the /home directory partition. Configure secret.doc (with a user and group owner of galleys) in /home/galleys with ACLs that allow user michael read-write access.

Set up appropriate partitions as a RAID 1 array (with one spare partition), dedicated to the /home/pcplane directory. While you could do this during the installation process, do so after installation, for the purpose of this exercise.

Server Installation Problem: RHCE-Level Skills

In this part of the exam, you'll configure a number of different servers on the RHEL system. Assume this computer is on a gateway between your LAN and an external network such as the Internet. Based on the configuration shown in

Table B-1, set up /var on an LVM array.

Set up both a regular and a secure Web server. Make sure the home pages for each server are different. Limit access to within the LAN only, and to the users avionics, seats, and galleys. Create and activate a Web proxy server. Configure a Samba server that allows users to access their home directories from remote computers on the LAN. Create an NFS share, with full privileges, on your /tmp directory. Make sure SELinux settings support access to this share.

Set up an FTP server that supports only anonymous access, even from outside your LAN. Configure sendmail to support access from within the LAN; do not require address confirmation from a DNS server. Configure an incoming e-mail service that supports regular, non-secure IMAP4 connections. Activate the Secure Shell service, and allow access from inside and outside the LAN. Do not allow direct root logins through the Secure Shell connection.

Edit the Kickstart file that is created; set it up to be usable for other computers with an identical hardware configuration. The Kickstart file should also support the creation of the same partitions.

Installation Discussion

Since you can set up a Red Hat Enterprise Linux configuration in several ways, there is no one right answer for the listed requirements. But you should remember a few general concepts. It's normally fastest to include packages during the installation process. It's easiest (and generally faster) to set up RAID arrays and LVM groups during the installation process. Make sure that the services you set up are active at the appropriate runlevels.

You can set up DHCP addressing through the Red Hat installation program or in /etc/sysconfig/network. You'll also want to allow incoming connections to your SSH and FTP servers. You can do this with the Security Level Configuration tool, commands in /etc/hosts.allow and /etc/hosts.deny, directives in service-level configuration files, or even with appropriate iptables commands.

You can connect to a local printer by editing the files in /etc/cups or using the Printer Configuration utility. You can limit networking to working hours using appropriate cron jobs, stored in the /etc/cron.daily directory. Setting up an NIS client means activating the ypbind daemon and using **domainname** to designate the biglan NIS domain, or you can use the Authentication Configuration tool. Also, activate the SELinux **allow_ypbind** (Allow Daemons To Run With NIS) setting. Before the automounter works, you have to activate the **autofs** service as well as the appropriate commands in /etc/auto.master and /etc/auto.misc.

To support IP forwarding, you'll need to set the **net.ipv4.ip_forward** variable in /etc/sysctl.conf and activate it in the /proc/sys/net/ipv4/ip_forward file. You can install the RPMs of your choice, including system-config-boot, with the appropriate **rpm -ivh** ***packagename*** command; if there are dependencies, and you're connected to an appropriate repository, you can use the **yum install** ***packagename*** command. When you upgrade the Linux kernel, however, you should install it with **rpm -i** just in case the new kernel doesn't work. When you set up users in a special directory, don't forget to set up the directory with the SGID bit.

To make the /home directory work with quotas and ACL, you'll need to add the **usrquota** and **acl** options to the associated directive in /etc/fstab. Before you can configure quotas, you'll need to remount /home with at least the **usrquota** and **acl** settings. To give user michael read-write permissions to secret.doc in /home/galleys, set appropriate permissions to /home/galleys:

```
#
# chmod 701 /home/galleys/
```

Glossary

As the Red Hat exams are an advanced challenge, I limit this glossary to what you would see beyond the prerequisites; don't expect to see most basic terms from Chapter 1 here.

A-C

Access Control Lists (ACLs)

Access Control Lists (ACLs) provide an additional layer of access control to files and directories; associated with the setfacl and getfacl commands.

Address Resolution Protocol (ARP)

A protocol that maps an IP address to the hardware address on a network card.

anacron

The anacron service is designed to run cron jobs that could not run while a server was powered down.

Apache Web server

The Apache Web server provides both normal and secure Web services, controlled by the **httpd** daemon.

apachectl

The apachectl command is the preferred method to start and stop an Apache server.

arp (Address Resolution Protocol)

The **arp** command is used to view or modify the kernel's ARP table. Using **arp**, you can detect problems such as duplicate addresses on the network. Alternatively, you can use **arp** to add the required entries from your LAN.

at

The at command is similar to cron, but it allows you to run a job on a one-time basis.

authentication

The way Linux checks the login rights of a user. Linux and Unix users are normally authenticated through use of a username and password, checked against /etc/passwd and related files.

automounter

The automounter can be configured to mount local and network directories on an as-needed basis. It's configured in /etc/auto.master, /etc/auto.misc, /etc/auto.smb, and /etc/auto.net.

BIND (Berkeley Internet Name Domain)

BIND is the Unix/Linux software that is used to set up a Domain Name System (DNS) service. The associated daemon is **named**.

BIOS

The BIOS is the Basic Input/Output System that runs basic commands when you power up your computer. The BIOS menu allows you to customize many options, including the sequence of boot media.

/boot

The directory with the main files required to boot Linux, including the Linux kernel and initial RAM disk. By default, /boot is mounted on a separate partition.

BOOTP

A TCP/IP protocol that sends IP address information from a remote DHCP server.

caching-only name server

A caching-only name server that performs many of the functions of a DNS server. It stores the IP address associated with recent name searches, for use by other computers on your LAN.

chage

The chage command manages the expiration date of a password.

chattr

The chattr command allows you to change file attributes.

chgrp

The chgrp command changes the group that owns a file.

chkconfig

The chkconfig command manages runlevel service information. It can activate or deactivate services. It can also

customize services at specific runlevels.

chmod

The chmod command changes the permissions on a file.

chown

The chown command changes ownership on a file.

CIFS (Common Internet File System)

CIFS is the Microsoft name for advances in its networking software. It's also covered by the latest version of Samba, 3.0, which is included with RHEL.

client

A client is a computer that accesses information or resources from a server.

CNAME (canonical name)

The CNAME is a way to assign several different names to a computer in a DNS database. For example, you can set up *www* as an alias for the computer with your Web server. CNAME records cannot be assigned to a mail server (MX) or a Start of Authority (SOA) record.

cron

A service that runs jobs on a periodic basis. It's configured in /etc/crontab; by default, it executes jobs in the /etc/cron.hourly, /etc/cron.daily, /etc/cron.weekly, and /etc/cron.monthly directories.

crontab

Individual users can run the crontab command to configure jobs that are run periodically.

CUPS (Common Unix Printing System)

CUPS is the default print service for RHEL.

D-E

daemon

A process such as the Web service (**httpd**) or X Font Server (**xfs**) that runs in the background and executes as required.

/dev

The directory with device files, used to represent hardware and software components.

DHCP (Dynamic Host Configuration Protocol)

DHCP clients lease IP addresses for a fixed period of time from a DHCP server on a local network. The BOOTP protocol allows DHCP clients to get IP address information from a remote DHCP server. The DHCP server daemon is **dhcpd**; the DHCP client daemon is **dhclient**.

Disk Druid

Anaconda's hard disk management program. While the functionality is similar to fdisk and **parted**, Disk Druid is easier to use. However, it is available only during the Linux installation process.

display manager

A Linux display manager includes a dialog box for your username and password. Two major display managers are used in RHEL: **gdm** (GNOME) and **kdm** (KDE).

dmesg

The dmesg command lists the kernel ring buffer and the initial boot messages. If your system successfully boots, /var/log/dmesg is one place to look for messages if you think you have boot problems.

DNS (Domain Name System)

The DNS service maintains a database of fully qualified domain names such as www.redhat.com and IP addresses such as 206.132.41.202. If the domain name is not in the local database, DNS is normally configured to look to other, more authoritative, DNS servers. The associated daemon is **named**.

Dovecot

The Dovecot service is associated with POP and IMAP e-mail.

dual-core / multi-core

A dual-core CPU is one type of multiple-core CPU in which one physical integrated circuit includes two or more CPUs.

dumpe2fs

The dumpe2fs command provides a lot of information about the format of a partition.

e2label

The e2label command associates a device with a label, typically a filesystem directory.

edquota

The edquota command edits the quota for a user or a group.

emacs

The emacs editor is a popular text editor that can be run from a text console.

environment

Each user's environment specifies default settings such as login prompts, terminals, the PATH, mail directories, and more.

/etc/fstab

The/etc/fstab configuration file defines default mounted directories.

/etc/inittab

The /etc/inittab configuration file sets the default runlevel and starts key processes such as terminal gettys.

/etc/X11/prefdm

The /etc/X11/prefdm configuration file specifies the preferred GUI display manager.

exportfs

The exportfs command allows shared NFS directories to be shared with a network.

ExpressCard

An ExpressCard is the successor to the PC Card/PCMCIA standard. The two standards are not compatible, and PCMCIA cards do not fit into ExpressCard slots.

F-H

fdisk

A standard disk partition command utility that allows you to modify the physical and logical disk partition layout.

Fedora Linux

The successor to the freely available version of Red Hat Linux; more information on this Linux distribution is available online at www. fedoraproject.org. Formerly known as Fedora Core Linux, starting with Fedora 7, it includes all former Core and Extras packages.

filesystem

Filesystem has multiple meanings in Linux. It refers to mounted directories; the root directory (/) filesystem is formatted on its own partition. It also refers to file formats; Linux partitions are typically formatted to the ext3 filesystem.

Filesystem Hierarchy Standard

The Filesystem Hierarchy Standard is the official way to organize files in Unix and Linux directories. The top-level directory is known as the root directory (/); users' home directories are configured in /home.

find

The find command searches for a desired file through a given directory and its subdirectories.

fips

The First Interactive Partition Splitter, fips, allows you to split existing VFAT partitions.

firewall

A hardware or software system that prevents unauthorized access over a network. Normally used to protect a private LAN from attacks through the Internet.

firstboot

The process that starts when you've configured RHEL during installation to boot into the GUI (runlevel 5); also known as First Boot.

fsck

The fsck command checks the filesystem on a Linux partition for consistency.

FTP (File Transfer Protocol)

The FTP protocol is a TCP/IP protocol designed to optimize file transfer between computers.

gateway

A gateway is a route from a computer to another network. A default gateway address is the IP address of a computer or router that connects a LAN with another network such as the Internet.

getfacl

The getfacl command lets you read Access Control Lists (ACLs) on files and directories.

getty

A getty is a terminal program that includes prompts for a login and a password. Virtual console gettys are configured through the **mingetty** program via /etc/inittab.

GNOME (GNU Network Object Model Environment)

GNOME is the default GUI desktop for Red Hat Enterprise Linux.

GPG (GNU Privacy Guard)

GPG is an implementation of the OpenPGP standard included with Red Hat Enterprise Linux.

group ID

Every Linux group has a group ID, as defined in /etc/group.

GRUB (Grand Unified Bootloader)

The default boot loader for RHEL.

grub-install

The grub-install command makes your BIOS look for your GRUB boot loader.

hard limits

Associated with user quotas. Specifies the permanent maximum amount of space a user can have on a partition, independent of grace periods.

home directory

The home directory is the login directory for Linux users. Normally, this is /home/*user*, where *user* is the user's login

name. It's also represented by the tilde (~) in any Linux command.

htpasswd

The htpasswd command helps create passwords for accessing a local Web site.

I-L

ICMP (Internet Control Message Protocol)

A protocol for sending online error control messages. Associated with the **ping** command.

ifconfig

The ifconfig command is used to configure and display network devices.

init

The init process is the first Linux process called by the kernel. This process starts other processes that compose a working Linux system, including the shell.

Initial RAM Disk

RHEL uses an initial RAM disk in the boot process; it's stored as an initrd-`uname -r`.img file in the /boot directory. You can create your own from the currently booted kernel with the **mkinitrd initrd-`uname -r`.img `uname -r`** command.

Internet Print Protocol (IPP)

The Internet Print Protocol (IPP) is the evolving standard for printers shared over networks. It's being adapted by all major operating systems; the Linux implementation is CUPS.

IP forwarding

IP forwarding is when data is forwarded between computers or networks through your computer.

iptables

The iptables command is the basic command for firewalls and masquerading.

IPv4, IPv6

IPv4 and IPv6 are different systems of IP addressing. Version 4 is what we use today and is based on 32-bit addresses; version 6 is coming on line and is based on 128-bit addresses.

iSCSI (Internet *SCSI)*

Internet SCSI is a network protocol standard, associated with SCSI-3 specifications on network storage devices.

KDE

A GUI for Linux and Unix computers. Also known as the K Desktop Environment.

Kdump

The Kdump service allows you to configure what happens in the event of a kernel crash. You can dedicate a specific amount of RAM to the process, which is then unavailable for other processes.

kernel

The kernel is the heart of any operating system. It loads device drivers. You can recompile a Linux kernel for additional drivers, for faster loading and to minimize the required memory.

kernel module

Kernel modules are pluggable drivers that can be loaded and unloaded into the kernel as needed. Some loaded kernel modules are shown with the **lsmod** command.

Kickstart

Kickstart is the Red Hat automated installation system that allows you to supply the answers required during the installation process. When properly configured, a kickstart floppy can allow you to start your computer and install RHEL automatically from a network source.

LDP (Linux Documentation Project)

The LDP is a global effort to produce reliable documentation for all aspects of the Linux operating system. Its work is available online at www.tldp.org.

lftp

The lftp command starts a slightly more flexible FTP command line client.

Lightweight Directory Access Protocol (LDAP)

The Lightweight Directory Access Protocol allows you to keep authentication information on a central server on your network.

locate

The locate command searches through a default database of files and directories. The database is refreshed daily with the mlocate.cron script in the /etc/ cron.daily/ directory.

logrotate

The logrotate command utility allows you to maintain log files. By default, RHEL uses the cron daemon to rotate, compress, and remove various log files.

lpc

You can use the lpc command to scan all configured print devices and queues.

lpq

You can use the lpq command to view print jobs still in progress.

lpr

You can use the lpr command to send print requests.

lprm

You can use the lprm command to remove print jobs from the queue.

logical extent (LE)

A logical extent (LE) chunk of disk space that corresponds to a physical extent (PE).

logical volume (LV)

A logical volume (LV) is composed of a group of logical extents (LEs).

Logical Volume Management (LVM)

Logical Volume Management (LVM) allows you to set up a filesystem on multiple partitions. Also known as the Logical Volume Manager.

lsattr

The lsattr command lists file attributes.

lvcreate

The lvcreate command creates a logical volume (LV) from a specified number of available physical extents (PEs).

lvdisplay

The lvdisplay command specifies current configuration information for logical volumes (LVs).

lvextend

The lvextend command allows you to increase the physical volume (PV) area allocated to a logical volume (LV).

lvremove

Functionally opposite to the lvcreate command.

M-P

masquerading

Masquerading enables you to provide Internet access to all of the computers on a LAN with a single public IP address.

MBR (Master Boot Record)

The first sector of a bootable disk. Once the BIOS cycle is complete, it looks for a pointer on the boot disk's MBR, which then looks at a boot loader configuration file such as grub.conf to see how to start an operating system.

mdadm

The mdadm command can help you view and configure RAID arrays.

mkbootdisk

The mkbootdisk command can create a boot disk, customized for your system.

mkfs

The mkfs command can help you format a newly configured partition. Variations are available including **mkfs.ext3**, which formats to the default ext3 filesystem.

modprobe

You can use the modprobe command to control device modules to be installed.

mount

You can use the mount command to specify mounted partitions, or attach local or network partitions to specified directories.

mount.cifs and umount.cifs

The **mount.cifs** and **umount.cifs** commands, when properly configured, allow regular users to mount directories shared over a Microsoft Windows network through Samba.

NAT (Network Address Translation)

NAT is a feature associated with firewall commands such as iptables, which connects computers inside your LAN to the Internet while disguising their true IP addresses. NAT modifies IP packet headers. The process is reversed for return messages. Closely related to masquerading.

netstat

The netstat command displays connectivity information for your network cards. For example, the **netstat -r** command is used to display the routing tables as stored in your kernel.

Network Time Protocol (NTP)

The Network Time Protocol allows you to synchronize your computer with a central timeserver. You can do this on RHEL with the Date/Time Configuration tool or by editing /etc/ntpd.conf and activating the ntpd service.

NFS (Network File System)

NFS is a file-sharing protocol originally developed by Sun Microsystems; it is the networked filesystem most commonly used for networks of Linux and Unix computers.

NIC (Network Interface Card)

A NIC connects your computer to a network. A NIC can be anything from a Gigabit Ethernet adapter to a telephone modem.

NIS (Network Information System)

NIS allows you to share one centrally managed authorization database for the Linux and Unix systems on your network.

PAM (Pluggable Authentication Module)

PAM separates the authentication process from individual applications. PAM consists of a set of dynamically loadable library modules that configures how an application verifies its users before allowing access.

parted parted

is a standard disk partition command utility that allows you to modify the physical and logical disk partition layout. Be careful when using it, as changes are immediately written to the partition table.

partprobe

You can use the partprobe command to reread a recently changed partition table without rebooting.

PATH

A shell variable that specifies the directories (and in what order) the shell automatically searches for input commands

and files.

PGP (Pretty Good Privacy)

A technique for encrypting messages, often used for e-mail. It includes a secure private- and public-key system similar to RSA. The Linux version of PGP is known as GPG (GNU Privacy Guard).

physical extent (PE)

A chunk of disk space created from a physical volume (PV) for Logical Volume Manager (LVM).

physical volume (PV)

An area of space for Logical Volume Manager (LVM) that usually corresponds to a partition or a hard drive.

Pirut

Pirut is the name of the RHEL package management tool.

Primary ATA (PATA)

Primary ATA is the media standard associated with older IDE drives, also known as ATA (Advanced Technology Attachment).

Primary Domain Controller (PDC)

A PDC is the governing server on a Microsoft Windows NT 4 network. You can configure RHEL with Samba to function as a PDC or as a member server on more current Microsoft networks.

/proc

/proc is the Linux *virtual* filesystem. *Virtual* means that it doesn't occupy real disk space. /proc files are used to provide information on kernel configuration and device status.

public/private key

Encryption standards such as PGP, GPG, or RSA are based on public/private key pairs. The private key is kept on the local computer; others can decrypt it with the public key.

Pup

Pup is short for the Package Updater, which monitors the Red Hat Network (RHN) for packages available for update.

pvcreate

The pvcreate command allows you to configure physical extents (PEs) from a properly configured partition.

pvdisplay

The pvdisplay command specifies current configuration information for physical volumes (PVs).

Q-R

quota
In Linux, a quota can limit users and/or groups by number of inodes or disk space. Quotas can include hard and soft limits.
quotacheck
The quotacheck command scans and creates user and group quota files.
quotaon
The quotaon command activates configured quotas.
RAID (Redundant Array of Independent Disks)
RHEL supports software RAID. You can use Anaconda to set up software RAID 0, 1, 5, and 6 arrays. You can also set up RAID arrays using the fdisk or **parted** command with mdadm. Also known as Redundant Array of Inexpensive Disks.
RAID 0
A RAID 0 array requires two or more partitions or hard drives. Reads and writes are done in parallel, increasing performance, filling up all partitions or hard drives equally. RAID 0 includes no redundancy; if any partition or hard drive in the array fails, all data in the array is lost.
RAID 1
A RAID 1 array requires two or more partitions or hard drives. RAID 1 is also known as mirroring, because the same information is written to both partitions. If one disk is damaged, all data will still be intact and accessible from the other disk.
RAID 5
A RAID 5 array requires three or more partitions. Parity information is striped across all partitions. If one partition fails, the data can be rebuilt. It can be automatically written to a spare disk.
RAID 6
A RAID 6 array requires four or more partitions. Parity information is striped twice across all partitions. If one or two partitions fail, the data can be rebuilt. It can be automatically written to a spare disk.
Red Hat Certified Engineer (RHCE)
Perhaps the elite certification available for Linux systems administrators. Designed to qualify Linux administrators with significant experience in configuring Linux LANs with Red Hat Enterprise Linux.
Red Hat Certified Technician (RHCT)
Another elite certification for newer Linux administrators. Designed to qualify Linux administrators with significant experience in configuring Linux workstations with RHEL. RHCEs must also meet all RHCT requirements.
Red Hat Hardware Compatibility List
The Red Hat Hardware Compatibility List (HCL) specifies all hardware that has been tested on systems running the various Red Hat operating systems. Red Hat provides installation support for any hardware that is listed as "support" on their HCL.
Red Hat Network (RHN)
The Red Hat Network (RHN) supports remote control and administration of systems with RHN subscriptions.
Red Hat Package Manager (RPM)
The Red Hat Package Manager is a system that sets up software in discrete packages. The associated **rpm** command allows you to add, remove, and upgrade packages.
refresh rate
The refresh rate regulates the rate at which the image you see on your screen is redrawn, in hertz (Hz).
repquota
The repquota command reports disk consumption.
resize2fs
The resize2fs command allows you to change the size of a filesystem, often used after increasing the space associated with an LVM.
reverse (inverse)zone
A DNS reverse (inverse) zone can be required by some servers, such as Apache and sendmail, to make sure an IP address points to a real computer. If the reverse zone host name does not match the IP address, the server might not

respond.

rndc

The rndc command is used to manage the operation of a DNS server; it's preferred over commands such as **service named start**.

root

This word has multiple meanings in Linux. The root user is the default administrative user. The root directory (/) is the top-level directory in Linux. The root user's home directory, /root, is a subdirectory of the root directory (/).

router

A computer that transfers messages between LANs. Computers that are connected to multiple networks often serve as routers.

rpmbuild

The rpmbuild command allows you to build source code based on information in a .spec file.

runlevel

RHEL includes six available runlevels, as defined in /etc/inittab. Key runlevels include 1, single-user mode; 3, text login; and 5, GUI login.

S

Samba

The Linux and Unix implementation of the Server Message Block protocol and the Common Internet File System (CIFS). Allows computers that run Linux and Unix to communicate with computers that run Microsoft Windows operating systems. I expect Samba 4.0, when released, to provide nearly full functionality as a Microsoft Active Directory Domain Controller.

secure virtual hosts

You can configure multiple secure virtual hosts on a single Apache server using the secure configuration file, /etc/httpd/conf.d/ssl.conf.

Security Enhanced Linux (SELinux)

An implementation of mandatory access control integrated into the Linux kernel; in essence, a different way of layering security within Linux.

sendmail

A standard e-mail server application used by most Internet e-mail.

Serial ATA (SATA)

The new standard on hard drives that makes it easier to chain hard drives in a series inside a physical system. SATA drives have device file labels similar to SCSI; for example, the first SATA drive is known as /dev/sda.

server

A computer that controls centralized resources such as files and printers. Servers can share these resources with client computers on a network.

setfacl

The setfacl command lets you control Access Control Lists (ACLs) on files and directories.

SGID

The SGID bit sets common group ID permissions on a file or directory.

Shadow Password Suite

The Shadow Password Suite creates an additional layer of protection for Linux users and groups in the /etc/shadow and /etc/gshadow files.

showmount

The showmount command lists the shared directories from an NFS server.

single-user mode

When you start RHEL in single-user mode, you're automatically logged in as the root user, without networking or most services. If your Linux system has boot problems, single-user mode may allow enough access to fix the problem.

smbpasswd

The smbpasswd command helps you create usernames and passwords for a Samba (Microsoft Windows) network.

SMTP (Simple Mail Transfer Protocol)

SMTP is a TCP/IP protocol for sending mail; used by sendmail.

SOA (Start of Authority)

In a DNS database, the SOA record is the preamble to all zone files. It describes the zone, the DNS server computer (such as ns.*your-domain*.com), the responsible administrator (such as hostmaster@your-domain.com), the serial number associated with this file, and other information related to caching and secondary DNS servers.

soft limit

Associated with user quotas. Specifies the maximum amount of space a user can have on a partition. Soft limits can be configured with grace periods.

spec file

Spec files are associated with source RPMs (SRPMS). You can modify an SRPM spec file to change the way an RPM package is built.

Squid

Squid is a high-performance HTTP and FTP caching proxy server.

SRPM (source RPM)

SRPMs include the source code required to build a binary RPM package. SRPMs are installed with the **rpm -i** command, which installs SRPM files within the /usr/src/redhat directory. You can then use the rpmbuild command to

create a binary RPM.

Structured Query Language (SQL)

The basis for several database systems that can be run on Linux, including MySQL and PostgreSQL.

SUID

The SUID bit sets common user ID permissions on a file or directory.

superuser

The superuser represents a regular user who has taken root user privileges. Closely associated with the **su** and **sudo** commands.

swap space

Linux uses swap space for less frequently used data that would otherwise be stored in RAM. It is normally configured in Linux in a swap partition.

system-config-*

Red Hat has created a series of GUI configuration tools to help configure a number of different systems and services. You can start them with a number of different commands that start with system-config-*. While it's usually faster to configure a configuration file directly, not every experienced administrator knows every detail of every major configuration file.

T-X

TCP/IP (Transmission Control Protocol/Internet Protocol)
TCP/IP is a suite of communications protocols for internetwork communication. It is primarily used as the communication system for the Internet.
Telnet
A terminal emulation program that allows you to connect to remote computers. RHEL includes the Kerberos version of the Telnet server, as configured through the /etc/xinetd.d/krb5-telnet configuration file.
tmpwatch
The tmpwatch command removes files that have not been accessed in a specified number of hours. The default daily tmpwatch script checks files in the /tmp and /var/tmp directories.
umask
The umask command defines default permissions for newly created files.
user ID (UID)
Every Linux user has a user ID, as defined in /etc/passwd.
usermod
The usermod command modifies different settings in /etc/passwd, such as expiration date and additional groups.
Very Secure FTP (vsFTP)
The Very Secure FTP service is the default FTP server for RHEL.
vgcreate
The vgcreate command creates a volume group (VG) from one or more physical volumes (PVs) for Logical Volume Manager (LVM).
vgdisplay
The vgdisplay command specifies current configuration information for volume groups (VGs).
vgextend
The vgextend command allows you to increase the extents or space allocated to a volume group (VG).
vi
The vi editor is a basic Linux text editor. While other editors are more popular, vi may be the only editor you have available in certain rescue environments.
virtual hosts
You can configure multiple Web sites on a single Apache server by configuring a number of virtual hosts in your /etc/httpd/conf/httpd.conf configuration file.
virtualization
Virtualization is an abstraction of computer resources; most often associated with platform virtualization, in which you can include one or more virtual machines on a physical system. Two options for virtualization are VMware and Xen.
VMware
VMware is a proprietary system with virtualization products freely available to all. With snapshots, it can help you test a system with less risk. I've written much of this book with RHEL installed on a VMware Server.
volume group (VG)
A collection of physical volumes (PVs) in Logical Volume Manager (LVM).
window manager
The window manager is a special type of X client that controls how other X clients appear on your display.
WINS (Windows Internet Name Service)
WINS provides name resolution on Microsoft networks; it can be activated on Samba.
X client
An X client is an application that uses the X server services to display output.
X Display
The X Display is a console and a virtual window. By default, there are six virtual text consoles configured with Linux; the X Display is associated with virtual console number seven.
X server
The X server is the part of the X Window System that runs on your desktop. The X server draws images on your screen, takes input from your keyboard and mouse, and controls access to your display.

X Window System

The GUI for Linux is also known as the X Window. Unlike other applications, the X Window System is a layered application.

Xen

Xen is the native virtualization technology to RHEL. It requires the use of a custom Xen kernel and can support virtual machines in paravirtualized and fully hardware virtualized modes.

xhost

The xhost command can be used to allow other hosts to access your X server. In other words, you can configure remote X clients to send their display to the local X server.

xinetd daemon

The **xinetd** "super-server" daemon controls connections to servers in the /etc/xinetd.d directory such as the **rsync** and Kerberos Telnet servers.

X.org

The X.org server is the default X server for RHEL.

Y

ypbind

The NIS client service is ypbind.

ypserv

The NIS server service is ypserv.

yum

The yum command allows you to update and install RPMs from remote sources, including dependencies. On RHEL 5, yum has replaced **up2date** for updates.

Index

Symbols

Index

A

Index

B

PREV NEXT

PREV NEXT

Index

C

Index

D

Index

E

Index

F

PREV NEXT

Index

G

Index

H

Index

I

Index

J

Index

K

Index

L

Index

M

Index

N

Index

O

Index

P

Index

Q

Index

R

Index

S

Index

T

Index

U

Index

V

Index

W

Index

X

Index

Y

Index

Z

List of Figures

Chapter 16: Troubleshooting

List of Tables

List of Exercises

List of Exam Details